Université de Montréal

The Normativity of Epistemic Rationality

par Marc-Kevin Daoust

Département de philosophie
Faculté des études supérieures et postdoctorales

Thèse présentée
en vue de l'obtention du grade de docteur
en philosophie

1er mai 2019

Université de Montréal

Département de philosophie, Faculté des études supérieures et postdoctorales

Cette thèse intitulée

The Normativity of Epistemic Rationality

Présentée par
Marc-Kevin Daoust

A été évaluée par un jury composé des personnes suivantes

Christine Tappolet
Présidente-rapporteuse

Daniel Laurier
Directeur de recherche

Aude Bandini
Membre du jury

Michael G. Titelbaum
Examinateur externe

Résumé

Cette thèse tente de démontrer que la rationalité épistémique est normative, ou plus précisément, que les agents doivent être épistémiquement rationnels. On peut dire provisoirement que la rationalité est un code – ou un ensemble de règles – en ce sens qu'elle correspond à une collection systématique d'exigences. En particulier, la rationalité est souvent identifiée à un ensemble systématique de lignes de conduite ayant pour but de faire en sorte qu'un agent réponde correctement à ses raisons, qu'il évite de se contredire, d'être incohérent ou akratique. La thèse poursuit aussi certains objectifs secondaires, en particulier : (i) montrer que les raisons épistémiques apparentes de croire que P (comprises comme des *propositions apparemment vraies* qui, si elles étaient vraies, militeraient en faveur de la conclusion que P) sont normatives ; (ii) montrer qu'il n'y a pas de dilemme insoluble de la rationalité épistémique ; (iii) montrer que, relativement à un ensemble de données probantes ou de raisons épistémiques, un agent idéalement rationnel n'a jamais la permission épistémique de croire que P et de croire que ~P simultanément. Si ces objectifs secondaires sont intéressants en eux-mêmes, ils contribuent aussi à confirmer l'idée selon laquelle la rationalité épistémique est normative.

Mots-clés : Philosophie, épistémologie, rationalité épistémique, normes épistémiques, données probantes, akrasie, cohérence, permissivisme épistémique

Abstract

This thesis argues that epistemic rationality is normative, or that agents ought to be epistemically rational. The property of rationality is here understood as a code. Specifically, the code of epistemic rationality requires various things, such as responding correctly to epistemic reasons one has, remaining coherent and avoiding akratic combinations of beliefs. Additionally, this thesis has secondary aims, such as: (i) arguing that apparent epistemic reasons to believe P (understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P) are deontically significant; (ii) arguing against unsolvable normative dilemmas of epistemic rationality; (iii) arguing against a specific type of permissiveness which roughly states that, relative to a body of epistemic reasons, it can be epistemically rational for an ideal agent to believe P and to disbelieve P. While these secondary aims are interesting in their own right, they confirm the main claim of this thesis, namely, that epistemic rationality is normative.

Keywords: Philosophy, epistemology, epistemic rationality, epistemic norms, evidence, akrasia, coherence, epistemic permissiveness

Table des matières

Liste des tableaux

Liste des thèses ou principes importants

Approximation Argument. Epistemic norms bear deontic significance because they are a sufficiently good approximation of epistemic ideals. However, responding to apparent reasons (and, especially, to non-factive apparent reasons) is not part of ideal epistemic scenarios. So, responding to apparent reasons does not bear deontic significance.

Consistency. Epistemic rationality requires that, if A believes that P, then it is false that A believes that ~P.

Datum. Epistemic reasons, understood in the narrow sense of *ordinary facts* counting in favour of the conclusion that P, are *deontically significant*. That is, an agent's epistemic permissions and obligations are affected by the balance of epistemic reasons he or she has.

Deontic Significance of Apparent Reasons. Apparent epistemic reasons (understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P) are *deontically significant*. That is, an agent's epistemic permissions and obligations are affected by the balance of apparent epistemic reasons he or she has.

Diachronic Prohibition. If agent A has a specific set of rational standards at time t_0 and does not acquire new epistemic reasons between t_0 and t_1, A should refrain from changing his or her rational standards at time t_1.

Elimination. Necessarily, if an agent takes inconsistent attitudes towards P at time t, then he or she violates a substantive requirement of epistemic rationality other than Consistency.

Epistemic Teleology Argument. Epistemic norms bear deontic significance because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). However, responding to apparent non-factive reasons is not necessarily truth-conducive. So, responding to apparent non-factive reasons does not necessarily bear deontic significance.

Intra-Level Coherence. Epistemic rationality requires that, if A believes that P_1, believes that P_2, ..., and believes that P_n, then it is false that A believes that $\sim(P_1 \wedge P_2 \ldots \wedge P_n)$.

Inter-Level Coherence (or Reasons Enkrasia). Epistemic rationality requires that, if A believes that he or she has sufficient epistemic reason to believe P, then A believes that P.

Modest Reasons-Responsiveness Thesis. Rationality consists (in part) in responding to reasons one has. No *impermissive* body of epistemic reasons supports both the beliefs that P and that ~P. However, in cases where both believing P and disbelieving P are warranted by a body of epistemic reasons, Consistency plays a distinct explanatory role.

Normativity:

Contributory Normativity. If epistemic rationality requires of A to believe P, then A has a reason to believe P.

Minimal Normative Hypothesis. Agents ought to respond correctly to (apparent) sufficient epistemic reasons they have. Responding correctly to (apparent) reasons one has is deontically significant.

Normativity+ of Consistency. Given the accuracy-dominance arguments, A ought to be consistent.

Normativity- of Consistency. Given the accuracy-dominance arguments, A ought not to be inconsistent.

Strict Constitutive Normativity. If epistemic rationality requires of A to believe P, then A ought to believe P *because* A is rationally required to believe P.

Strict Normativity. If epistemic rationality requires of A to believe P, then A ought to believe P.

Owned Apparent Reasons Thesis. Epistemic rationality requires that, if A has sufficient *apparent* epistemic reason to believe P, A believes that P. Apparent epistemic reasons are here understood as *apparent* facts, which are not necessarily facts.

Owned Reasons Thesis. Epistemic rationality requires that, if A has sufficient epistemic reason to believe P, A believes that P. Epistemic reasons are here understood as facts. If it appears to A that some propositions are true while they are false, A is not rationally required to respond to such apparent reasons.

Permissiveness:

Extreme Reasons Permissiveness. Relative to a body of epistemic reasons, A can be rationally permitted to believe P and to believe ~P.

Interpersonal Permissiveness. Relative to a body of epistemic reasons, two agents can be rationally permitted to entertain distinct incompatible beliefs. For example, agent 1 could be rationally permitted to believe P and agent 2 could be rationally permitted to believe ~P.

Permissive Epistemic Standards Thesis. The weight of epistemic reasons in favour of P is not objectively determined—rather, in order to determine their weight, epistemic reasons have to be subjectively mediated through an agent's rational epistemic standards. Furthermore, there are multiple incompatible rational epistemic standards, in the sense that there are incompatible but equally rational ways to draw various conclusions from a body of epistemic reasons.

Rational Closure (under conjunction). Necessarily, if (i) A is epistemically rational in believing P, (ii) A is epistemically rational in believing Q, and (iii) A comes to believe (P^Q) through a good reasoning process (such as deductive reasoning), then A is epistemically rational in believing (P^Q).

Reductionism:

Modest Reductionist Hypothesis. *In ideal theories of epistemic rationality*, putative requirements such as Inter-Level Coherence, Intra-Level Coherence or Consistency

have no distinct explanatory role when compared with substantive requirements of epistemic rationality.

Robust Reductionism. Putative structural requirements of epistemic rationality can be explained in terms of failures to satisfy substantive requirements of epistemic rationality.

Reliability Criterion. In the right conditions, if A is ideally rational, then A satisfies some available epistemic standards that optimize his or her ratio of true to false beliefs (and such standards lead A to reach the right answer more than 50% of the time).

Strong Reasons-Responsiveness Thesis. Epistemic rationality consists in responding to epistemic reasons one has. No body of epistemic reasons can simultaneously support the belief that P and the belief that ~P. So, requirements such as Consistency have no distinct explanatory role.

Superior Standard Thesis. Let $\{St_1, St_2, ..., St_n\}$ be a set of incompatible available epistemic standards containing n elements. Let Pr_i be the probability that satisfying St_i will lead one to form a true belief. Finally, assume that such standards are equally reliable (such that $Pr_1=Pr_2=...Pr_n$ and $0.5<Pr_i<1$). Then, there exists an available epistemic standard that is more reliable than St_1, St_2, ..., and St_n.

À ma mère

Remerciements

Mes remerciements vont en premier lieu à mon directeur, Daniel Laurier, qui m'a offert un appui extraordinaire pendant toutes ces années. J'ai été initié à la constellations de problèmes étudiés dans cette thèse dans les séminaires de Daniel. Il a été un directeur de thèse absolument exceptionnel tant sur le plan intellectuel qu'humain. Tous les arguments de cette thèse ont été améliorés et bonifiés suite à des discussions approfondies avec lui. Il faut aussi mentionner que, lorsque j'ai commencé mon doctorat en 2014, Daniel faisait partie des rares personnes, avec Christian Nadeau, qui croyaient sincèrement en mon potentiel à compléter une thèse de doctorat en philosophie. Je n'aurais jamais eu le courage d'entreprendre cette aventure sans leur appui. Je les remercie de tout mon cœur.

J'aimerais remercier ma famille et Louis-Xavier Proulx, mon conjoint, pour leur soutien inconditionnel pendant mon passage au doctorat. Je tiens aussi à remercier Samuel Montplaisir, Fannie Achard, Sophia Rousseau-Mermans, Félix-Antoine Gélineau, Rémi Tison, David Montminy, Erika Olivaux et Aude Bandini, mes plus chers collègues du département de philosophie, avec qui j'ai eu la chance d'échanger tout au long de mon doctorat. Votre présence au département, vos commentaires et vos encouragements m'ont donné l'énergie de terminer cette aventure. Un merci tout spécial à Simon-Pierre Chevarie-Cossette et Samuel Dishaw qui, malgré la distance nous séparant, ont très souvent accepté de relire mes manuscrits et de me fournir de précieux commentaires sur mes arguments.

Lors de mes études doctorales, j'ai eu la chance de visiter l'Université d'Uppsala ainsi que l'Université Harvard pour approfondir mes recherches et rencontrer d'autres étudiants travaillant sur la rationalité épistémique. Andrew Reisner m'a accueilli à Uppsala. Il a passé en revue quatre de mes articles qui forment désormais l'essentiel des chapitres 3 à 5 de ma thèse. Selim Berker m'a accueilli à Harvard et dans son séminaire d'épistémologie. Il m'a offert des commentaires précieux sur le second chapitre de ma thèse. Je les remercie chaleureusement de m'avoir offert ces occasions inespérées.

J'aimerais remercier toutes les autres personnes ayant accepté de me donner des commentaires sur certains de mes arguments, soit Karl Bergman, Pat Bondy, Étienne Brown,

François Claveau, Charles Côté-Bouchard, Peter Dietsch, Joshua DiPaolo, Lidal Dror, Megan Entwistle, Caitlin Fitchett, Branden Fitelson, Daniel Fogal, Jens Gillessen, David Godden, Bruno Guindon, Ulf Hlobil, Sophie Horowitz, Tim Juvshik, Molly Kao, Hilary Kornblith, Matthew Kopec, Éliot Litalien, Jocelyn Maclure, Carl Montan, Molly O'Rourke-Friel, Michele Palmira, Richard Pettigrew, Andrei Poama, Olle Risberg, Juliette Roussin, Henrik Rydéhn, Matthew Scarfone, Xander Selene, Justin Snedegar, Quentin Soussen, Julia Staffel, Robert Stephens, Folke Tersman, Mike Titelbaum et Alex Worsnip. J'oublie sans doute certaines personnes et je les prie à l'avance de m'excuser. De plus, un merci à ces évaluateurs de l'ombre qui, lorsqu'ils sont de bonne humeur, acceptent de publier mes articles et m'offrent des conseils pertinents pour les améliorer.

Certains chapitres de cette thèse sont basés sur des articles publiés dans différentes revues. Le chapitre 3 est basé sur l'article « Epistemic Akrasia and Epistemic Reasons », publié dans *Episteme*. Le chapitre 4 est basé sur deux articles, dont « The Explanatory Role of Consistency Requirements » publié dans *Synthese*. Le chapitre 5 est basé sur plusieurs articles, dont « Epistemic Uniqueness and the Practical Relevance of Epistemic Practices » publié dans *Philosophia*, « One Standard To Rule Them All? » publié dans *Ratio*, ainsi que « Peer Disagreement and the Bridge Principle » publié dans *Topoi*. Je remercie ces revues d'autoriser la réutilisation du contenu de mes articles pour ma thèse.

Cette recherche a reçu l'appui financier du Groupe de recherche interuniversitaire sur la normativité (GRIN), du Conseil de recherches en sciences humaines du Canada (bourse #767-2016-1771) et du Fonds de recherche du Québec – société et culture.

Introduction

Projet principal de la thèse

Cette thèse tente de démontrer que la rationalité épistémique est normative, ou plus précisément, que les agents doivent être épistémiquement rationnels. Déjà, il n'est pas simple d'indiquer en quoi consiste la rationalité épistémique. Disons provisoirement que la rationalité est pensée comme un code – ou un ensemble de règles – en ce sens qu'elle correspond à une collection systématique d'exigences. En particulier, la rationalité est souvent identifiée à un ensemble systématique de lignes de conduite ayant pour but de faire en sorte qu'un agent réponde correctement à ses raisons, qu'il évite de se contredire, d'être incohérent ou akratique.

Par opposition aux normes de politesse ou d'étiquette, les normes issues de la rationalité épistémique semblent posséder une réelle autorité normative. Il y a un sens dans lequel on peut dire que la politesse ou l'étiquette sont « normatives », puisqu'il existe un ensemble de règles à respecter pour avoir de bonnes manières. Néanmoins, les exigences de l'étiquette ou la politesse ne sont pas normatives au sens fort du terme, en ce sens qu'une personne peut être dans *l'obligation* d'être impolie, qu'un agent peut être *louable* d'être impoli, qu'il peut y avoir une *valeur* associée à l'impolitesse dans certaines situations, et ainsi de suite. Par comparaison, la rationalité épistémique semble intimement associée à nos obligations, au blâme, à la valeur d'avoir certaines croyances, et ainsi de suite.

La question de la normativité de la rationalité épistémique peut sembler banale. Or, il n'en est rien : des problèmes complexes se cachent derrière les apparences. Une des sources de ces problèmes est que nous voulons généralement deux choses: des croyances vraies et des croyances rationnelles. Mais il se peut qu'une croyance rationnelle soit fausse. Alors, ces desiderata peuvent entrer en conflit.

Prenons le cas où un agent (appelons-le Kim) a d'excellentes raisons de croire une proposition fausse. Spécifiquement, l'hôtel de Kim est en feu, mais Kim a d'excellentes raisons de croire que son hôtel n'est pas en feu. Kim pourrait alors croire que la meilleure chose à faire est de rester dans sa chambre. Cette tension entre croyance vraie et croyance

rationnelle aura des répercussions pratiques évidentes : il serait préférable pour Kim de se diriger vers l'escalier de secours, mais comme il est rationnel de croire qu'il devrait rester dans sa chambre, il sera sans doute peu motivé à se diriger vers l'escalier de secours.

Supposons que, dans la situation précédente, Kim croit qu'il devrait demeurer dans sa chambre. Devrait-on dire que Kim croit la bonne chose? Devrait-on dire que Kim est autorisé à croire une telle chose? Est-il blâmable s'il se met à croire qu'il devrait quitter sa chambre? Supposons que Kim cesse de croire qu'il devrait demeurer dans sa chambre et se met à croire qu'il devrait se diriger vers la sortie de secours. Est-ce une bonne chose que l'agent ait changé sa croyance? Kim est-il autorisé à changer sa croyance? Est-il blâmable d'avoir formé cette croyance?

Prenons un cas encore plus simple. Supposons que, à t_0, Kim croit que P et croit aussi que ~P. Kim est irrationnel, puisque ses croyances sont contradictoires. Maintenant, supposons que Kim cesse de croire que P à t_1. Ainsi, il cesse d'être irrationnel (du moins, à t_1, sa croyance que ~P n'est plus en contradiction avec ses autres croyances). Or, supposons que P est vrai. Ainsi, à t_0, Kim enfreint l'exigence de non-contradiction, puisqu'il a une croyance vraie et une croyance fausse. À t_1, Kim n'enfreint plus l'exigence de non-contradiction, mais il n'a plus de croyance vraie et il a toujours une croyance fausse. La question qui se pose est la suivante : en quoi Kim a-t-il progressé entre t_0 et t_1? Si la rationalité épistémique est normative, il semble que Kim aurait dû progresser entre t_0 et t_1, puisqu'il n'enfreint plus une exigence de rationalité (l'exigence de non-contradiction). Or, il semble que Kim n'a fait aucun progrès : pire, puisqu'il a abandonné une croyance vraie entre t_0 et t_1, il semble que Kim est dans une moins bonne position à t_1. *Au moins, à t_0, Kim avait une croyance vraie.*

Déjà, ces deux illustrations soulèvent bon nombre de questions que nous tenterons de résoudre au fil de cette thèse. Premièrement, de quoi parle-t-on lorsque l'on fait référence à la normativité de la rationalité épistémique? À quelles normes (ou à quelle normativité) fait-on référence? Et que savons-nous à propos de la rationalité épistémique? Par exemple, quels rôles jouent la cohérence et les raisons dans les théories de la rationalité? Deuxièmement, quels sont les problèmes auxquels font face les tenants de la normativité de la rationalité épistémique? Troisièmement, quelles sont les conditions nécessaires ou suffisantes pour répondre à ces

questions et offrir une justification convaincante de la normativité de la rationalité épistémique? Toutes ces questions seront analysées en temps et lieu.

Objectifs secondaires de la thèse

Dans cette thèse, je tente aussi de résoudre certaines questions périphériques à la normativité de la rationalité épistémique ayant fait l'objet d'intenses débats dans les dernières années. Voici un bref aperçu des questions périphériques qui m'intéresseront.

La normativité des raisons apparentes. Dans les quinze dernières années, de nombreux épistémologues ont tenté de déterminer les sources ultimes de la normativité épistémique. Une vaste majorité d'entre eux s'entendent pour dire que les raisons sont irréductiblement normatives. Les raisons de croire que P sont ici comprises comme des faits ordinaires qui militent en faveur de la conclusion que P. Par exemple, le fait que le ciel s'assombrit est une raison de croire qu'il va pleuvoir. Si j'ai cette raison (p. ex., si j'ai accès au fait que le ciel s'assombrit), j'ai une raison de croire qu'il va pleuvoir. Cette raison (et les autres que je possède en faveur de la conclusion qu'il va pleuvoir) *affecte* mon obligation épistémique (ou ma permission) de croire qu'il va pleuvoir.

Maintenant, qu'en est-il des raisons apparentes? Les raisons apparentes de croire que P sont ici comprises comme des propositions apparemment vraies qui, s'ils étaient vrais, militeraient en faveur de la conclusion que P. Par exemple, supposons que le ciel *m'apparait s'assombrir*, alors qu'en réalité, le ciel ne s'assombrit pas. La simple apparence que le ciel s'assombrit affecte-t-elle mon obligation (ou ma permission) de croire qu'il va pleuvoir? Les internalistes à propos des normes épistémiques croient que oui, alors que les externalistes à propos des normes épistémiques pensent que les raisons apparentes ne sont pas forcément normatives (ou même, pas du tout normatives). Dans le chapitre 2, de cette thèse, je tente d'élucider cette question.

Les conflits potentiels de la rationalité épistémique. Dans les dernières années, la question de *l'unité* ou de la *compatibilité* des différentes normes épistémiques a fait couler beaucoup d'encre. Par exemple, les exigences de la rationalité épistémique admettent-elles des

dilemmes normatifs insolubles? Par exemple, que se passe-t-il si l'exigence de rationalité A entre en conflit avec l'exigence de rationalité B? Supposons que, pour bien répondre à mes raisons épistémiques, je dois croire X et je dois aussi croire Y. Cependant, on peut aussi supposer que, afin d'être cohérent, il faut éviter de croire X et de croire Y simultanément. À supposer que les exigences de rationalité incluent des exigences touchant les raisons épistémiques ainsi que des exigences de cohérence, les exigences ci-dessus seront incompatibles entre elles. La question qui se pose alors est la suivante : à supposer qu'il existe de tels conflits, sont-ils solubles? Le chapitre 3 de ma thèse tente de répondre à cette question.

Le permissivisme épistémique. La sphère pratique admet bon nombre de situations permissives. Pensons à l'âne de Buridan. Imaginons qu'un âne affamé soit à égale distance de deux bottes de foin identiques. L'âne se trouve alors dans une situation *permissive* : il est rationnel pour l'âne d'avoir l'intention de manger la botte de foin à sa gauche, et il peut être rationnel pour l'âne d'avoir l'intention de manger la botte de foin à sa droite, même si ces options sont incompatibles entre elles.

La question de savoir si la sphère épistémique admet des situations permissives est bien plus complexe. Par exemple, il est naturel de penser que, si un agent a autant de raisons épistémiques de croire P que de croire ~P, il n'a ni la permission de croire P, ni la permission de croire ~P. L'agent devrait plutôt suspendre son jugement quant à P. Cette intuition persistante a mené bon nombre de philosophes à conclure que la rationalité épistémique est impermissive. En d'autres termes, relativement à un ensemble de données probantes ou de raisons épistémiques, un agent n'a jamais la permission épistémique de croire que P et de croire que ~P simultanément (ou de croire que P et de suspendre son jugement quant à P simultanément). Les chapitre 4 et 5 approfondissent ce débat.

Comment lire cette thèse

Ces débats périphériques ci-dessus sont en lien direct avec la question de la normativité de la rationalité épistémique. En effet, ma défense de la normativité de la rationalité épistémique repose en partie sur les réponses que j'offre à ces différents problèmes. Cela dit,

ces questions sont aussi intéressantes en elles-mêmes, indépendamment de leurs implications dans le débat sur la normativité de la rationalité épistémique.

Étant donné ce qui précède, j'ai choisi d'écrire ma thèse de la manière suivante : chaque chapitre peut être lu indépendamment des autres et tente de résoudre une des questions périphériques mentionnées précédemment. Cela dit, la conclusion de chaque chapitre fait le pont entre le débat entourant la normativité de la rationalité et les thèses défendues dans le chapitre. À titre d'exemple, le chapitre 2 défend l'idée selon laquelle les raisons apparentes sont normatives, et la conclusion du chapitre 2 explore ce que la normativité des raisons apparentes nous apprend sur la normativité de la rationalité épistémique.

Résumé des chapitres

Dans le chapitre 1, je fais l'état des lieux des différents débats entourant la normativité de la rationalité épistémique. Je propose deux hypothèses qui guideront l'ensemble de ma thèse. Ces hypothèses sont :

Minimal Normative Hypothesis. Les agents doivent répondre correctement à leurs raisons (apparentes) suffisantes. En d'autres termes, les raisons (apparentes) que l'on possède sont normatives.

Modest Reductionist Hypothesis. Dans les théories de la rationalité épistémique idéale, les exigences structurales de cohérence ne jouent aucun rôle explicatif distinct par rapport aux exigences substantielles de la rationalité épistémique.

J'explique ensuite que, si ces deux hypothèses sont confirmées, on pourra en conclure que la rationalité épistémique est normative. J'explique aussi comment les chapitres 2 à 5 contribueront à confirmer ces deux hypothèses.

Dans le chapitre 2, je soutiens la thèse selon laquelle les raisons apparentes sont normatives. Ce chapitre répond à deux objections courantes à l'endroit de la normativité des raisons apparentes. Selon la première objection, les normes épistémiques sont une approximation satisfaisante des idéaux épistémiques. Or, les raisons apparentes (en particulier, les raisons apparentes qui ne sont pas des faits) sont absentes des idéaux épistémiques. Selon la seconde objection, les normes épistémiques ont à voir avec la maximisation de la valeur épistémique. Par exemple, à supposer que la vérité soit un bien épistémique, les normes

épistémiques devraient tendre à optimiser le ratio de croyances vraies à fausses des agents.s biens. Or, répondre correctement aux raisons apparentes que l'on a (en particulier, aux raisons apparentes qui ne sont pas des faits) ne garantit pas un bon ratio de croyances vraies à fausses. Ces deux objections remettent sérieusement en cause l'idée selon laquelle les raisons apparentes sont normatives. Je réponds à ces critiques et je soutiens la thèse selon laquelle les raisons apparentes sont normatives en explorant un nouveau modèle des normes épistémiques. La particularité de ce modèle est qu'il prend en compte les contraintes légitimes sur les options épistémiques qu'ont les agents.

Dans le chapitre 3, j'explore la possibilité d'un conflit entre cohérence et réponse adéquate aux raisons épistémiques. Selon plusieurs auteurs, il existe des situations où un agent doit choisir entre demeurer cohérent et répondre correctement à ses raisons épistémiques. Dans ce contexte, et à supposer que les agents ne se trouvent jamais dans des dilemmes épistémiques insolubles, ou bien (i) la rationalité épistémique permet parfois aux agents d'être incohérents, ou alors (ii) la rationalité épistémique permet parfois aux agents de ne pas répondre correctement à leurs raisons épistémiques. À première vue, ces deux possibilités sont aussi insatisfaisantes l'une que l'autre. Je compare ce nouveau conflit à des paradoxes bien connus de la rationalité épistémique, notamment le paradoxe de la loterie. De plus, j'avance l'idée selon laquelle ces conflits sont solubles.

Dans le chapitre 4, j'étudie de près l'exigence de non-contradiction. Selon cette exigence, la rationalité exige que, si A croit que P, alors il est faux que A croit que ~P. La question est de savoir si cette exigence joue un rôle explicatif distinct dans la théorie de la rationalité épistémique. Plusieurs philosophes nient le rôle explicatif distinct de l'exigence de non-contradiction. Ils pensent notamment que, si un agent répond correctement à ses raisons épistémiques, il ne peut croire que P et croire que ~P simultanément (puisqu'aucun ensemble de raisons épistémiques ne soutient simultanément ces deux conclusions). Si ces philosophes ont raison, l'exigence n'a aucun potentiel explicatif distinct dans la théorie de la rationalité épistémique. Je défends une alternative à ce réductionnisme concernant l'exigence de non-contradiction. Plus précisément, j'avance la thèse selon laquelle l'exigence de non-contradiction joue un rôle explicatif dans les situations épistémiquement permissives, où les

raisons épistémiques d'un agent soutiennent à la fois la conclusion que P ainsi que la conclusion que ~P. Cette forme de permissivisme est une possibilité à prendre au sérieux lorsqu'il est question des standards épistémiques d'un agent.

Dans le chapitre 5, j'offre un nouvel argument contre le permissivisme des standards épistémiques *dans les théories idéales de la rationalité*. Mon argument repose en grande partie sur le théorème du Jury de Condorcet. De plus, j'explique pourquoi ce résultat pourrait ne pas être généralisable aux théories non idéales de la rationalité épistémique. Enfin, je réponds à plusieurs objections possibles contre mon argument.

Les conclusions de chaque chapitre (et la conclusion générale) exposent comment les arguments élaborés dans cette thèse militent en faveur de la normativité de la rationalité épistémique.

Introduction (English Version)

Main Project of the Thesis

This thesis tries to show that epistemic rationality is normative, or more specifically, that agents ought to be epistemically rational. Let us say tentatively that rationality is a code or a set of rules. Some requirements of epistemic rationality govern how an agent responds to his or her reasons. Other requirements govern combinations of beliefs, and aim at avoiding inconsistent or incoherent combinations of beliefs.

By way of contrast with norms of politeness or etiquette, the norms of epistemic rationality seem genuinely normative. There is a weak sense in which one can say that politeness or etiquette are "normative", since one ought to satisfy a set of rules in order to be polite. Nevertheless, the requirements of etiquette or politeness are not normative in a genuine sense, since an agent can be *required* to be rude. Relatedly, it can be commendable to be rude, and there can be value associated with rudeness in some contexts. By way of contrast, epistemic rationality seems intimately tied to our obligations as epistemic agents.

The debate on the normativity of epistemic rationality is complex. Here is an important problem with the normativity of epistemic rationality. From an epistemic point of view, we want two things: true beliefs and rational beliefs. But some rational beliefs are false. So, these desiderata can conflict with each other.

Suppose an agent (call him Kim) has excellent reasons to believe a false proposition. Specifically, Kim's hotel is on fire, but Kim has excellent reasons to believe that his hotel is not on fire. Kim could then believe that the best thing he can do is to stay in his room. His rational (but false) belief has clear practical implications. It would be better for Kim to get out of his room, but since it is reasonable for him to believe that he should stay in his room, he will probably lack motivation to get out of his room.

Now, suppose Kim believes that he should stay in his room. Should we say that Kim has a *good* belief? Should we say that Kim is *permitted* to believe such a thing? Is he

praiseworthy for having this belief? If he starts believing that he should leave his room, is he blameworthy? Is he permitted to revise his belief?

Consider a simpler case. Suppose that at t_0, Kim believes that P and also believes that ~P. Kim is irrational, since his beliefs are jointly inconsistent. Now, suppose that Kim stops believing that P at t_1. So, he stops being irrational (at least, at t_1, his belief that ~P is no longer inconsistent with his other beliefs). Assume that P is true. Thus, at t_0, Kim violates the requirement of consistency, and he has one true belief and one false belief. At t_1, Kim no longer violates the requirement of consistency, but he no longer has a true belief (and he still has a false belief). The question is: why should we think that Kim did any progress between t_0 and t_1? If epistemic rationality is normative, it seems that Kim should have made some progress between t_0 and t_1, since he no longer violates a requirement of rationality (the requirement of consistency). But it seems that Kim has made no progress. In fact, Kim seems worse off at t_1. *At least at t_0, Kim had a true belief.*

These examples raise many questions that I will try to answer in this thesis. First, when talking about the normativity of epistemic rationality, what are we talking about exactly? Which normativity and which rationality are we concerned with? What do we know about epistemic rationality? For example, what roles do coherence requirements and reasons-responsiveness requirements play in theories of rationality? Which challenges do proponents of the normativity of epistemic rationality face? What are the necessary or sufficient conditions to offer a satisfactory vindication of the normativity of epistemic rationality?

Secondary Objectives of the Thesis

In this thesis, I will also try to solve various peripheral issues concerning epistemic norms. While these issues are interesting in their own right, they confirm the main claim of this thesis, namely, that epistemic rationality is normative. Here is a brief overview of the issues I will try to solve.

The normativity of apparent reasons. In the past fifteen years, many epistemologists have tried to identify the ultimate sources of epistemic normativity. Plausibly, reasons are

irreducibly normative. Reasons for believing that P are here understood as ordinary facts (or true propositions) that count in favour of the conclusion that P. For example, the fact that the sky is getting darker is a reason to believe that it will rain. If I have this reason (e.g., if I have access to the fact that the sky is getting darker), I have a reason to believe that it will rain. This reason *affects* my epistemic obligation (or my epistemic permission) to believe that it will rain.

Now, what about the apparent reasons? Apparent reasons for believing that P are here understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P. For example, suppose it appears to me that the sky is getting darker, while the sky is, in fact, not getting darker. Does the mere appearance that the sky is getting darker affect my obligation (or my permission) to believe that it will rain? Internalists about epistemic norms say yes, while externalists about the epistemic norms say no. In chapter 2, I try to elucidate this question.

The conflicts of epistemic rationality. Recently, the question of the *unity* of epistemic norms has been subject to controversy. For example, are there unsolvable normative dilemmas of epistemic rationality? Can some requirements of epistemic rationality conflict with each other? Suppose that, in order to respond correctly to my epistemic reasons, I ought to believe that X and I also ought to believe that Y. However, assume that, in order to be coherent, I ought to avoid believing X and believe Y simultaneously. If epistemic rationality has to do with reasons-responsiveness requirements as well as coherence requirements, the above requirements are inconsistent with each other and indicate that rationality is *conflictual*. This doesn't square well with plausible assumptions concerning epistemic rationality. Chapter 3 sheds light on this puzzle.

Epistemic permissiveness (or permissivism). In the practical realm, there are permissive situations—think of Buridan's ass who faces equally good stacks of hay. Whether there are permissive situations in the epistemic realm is a much more complex issue. For example, it is natural to think that, if an agent has as equally good reason to believe P and to believe ~P, he or she should neither believe P nor disbelieve P. Rather, the agent ought to

suspend judgment of whether P. This intuition has led many philosophers to conclude that epistemic rationality is impermissive. In other words, relative to a body of evidence, an agent is never epistemically permitted to believe that P and believe that ~P simultaneously. Chapters 4 and 5 investigate this debate.

How to Read this Thesis

There is a close connection between my argument for the normativity of epistemic rationality and the secondary issues mentioned above. Yet these secondary issues are interesting in their own right. In view of the foregoing, I chose to write my thesis as follows: each chapter can be read independently of the others and tries to solve one of the secondary issues mentioned previously. However, the conclusion of each chapter explains how these secondary issues are connected to the debate on the normativity of rationality. For example, chapter 2 argues that apparent reasons are normative. In the chapter's conclusion, I explain how the normativity of apparent reasons is related to the normativity of epistemic rationality.

Summary of the Chapters

In chapter 1, I set the stage. After having introduced some essential notions and distinctions surrounding the notion of rationality, I present some important worries and complications in the debate about the normativity of epistemic rationality. Then, I introduce two hypotheses (the Minimal Normative Hypothesis and the Modest Reductionist Hypothesis) that are jointly sufficient for vindicating the normativity of epistemic rationality:

Minimal Normative Hypothesis. Agents ought to respond correctly to (apparent) sufficient epistemic reasons they have. Responding correctly to (apparent) reasons one has is deontically significant.

Modest Reductionist Hypothesis. In ideal theories of epistemic rationality, putative structural requirements such as Inter-Level Coherence, Intra-Level Coherence or Consistency have no distinct explanatory role when compared with substantive requirements of epistemic rationality.

That is, if the Minimal Normative Hypothesis and the Modest Reductionist Hypothesis are correct, epistemic rationality is deontically significant. At the end of the chapter, I explain how

we can vindicate both hypotheses and I state some basic assumptions that I will make throughout the thesis.

In chapter 2, I argue that apparent reasons are normative. That is, an agent's epistemic permissions and obligations are affected by the balance of apparent epistemic reasons he or she has. The chapter also responds to two objections against the normativity of apparent reasons. According to the first objection, epistemic norms bear deontic significance because they are a sufficiently good approximation of epistemic ideals. However, responding to apparent reasons (and, especially, to non-factive apparent reasons) is not part of ideal epistemic scenarios. According to the second objection, epistemic norms bear deontic significance because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). However, responding to apparent non-factive reasons is not truth-conducive. These two objections entail that responding to apparent non-factive reasons does not bear deontic significance. The normativity of apparent reasons and both objections are considered in a new framework (i.e. second-best epistemology), which takes into account the legitimate constraints on available belief-forming processes.

In chapter 3, I analyze a putative conflict between coherence requirements and reasons-responsiveness requirements. It seems that epistemically rational agents should avoid incoherent combinations of beliefs and should respond correctly to their epistemic reasons. However, some situations seem to indicate that such requirements cannot be simultaneously satisfied. In such contexts, assuming that there is no unsolvable dilemma of epistemic rationality, either (i) it could be rational that one's higher-order attitudes do not align with one's first-order attitudes or (ii) requirements such as responding correctly to epistemic reasons that agents have are not genuine rationality requirements. This result doesn't square well with plausible theoretical assumptions concerning epistemic rationality. So, how do we solve this puzzle? I suggest that an agent can always reason from infallible higher-order reasons. This provides a partial solution to the above puzzle.

In chapter 4, I shed light on the explanatory role and the normativity of the requirement of consistency. Is inconsistency a mere symptom of having violated other requirements of

rationality—notably, reasons-responsiveness requirements? Or is inconsistency irrational on its own? This question has important implications for the debate on the normativity of epistemic rationality. I defend a new account of the explanatory role of the requirement of consistency. Roughly, I will argue that, in cases where an epistemically rational agent is permitted to believe P and also permitted to disbelieve P (relative to a body of epistemic reasons), the consistency requirement plays a distinct explanatory role. I will also argue that such a type of permissiveness is a live possibility when it comes to rational epistemic standards.

In chapter 5, I argue that a specific type of permissiveness concerning epistemic standards is false with respect to ideal theories of epistemic rationality. Specifically, those who argue that epistemic standards are permissive can't make sense of the reliability criterion (at least in ideal theories of epistemic rationality). My strategy relies on Condorcet's Jury Theorem. I then explain why my strategy might not generalize to non-ideal theories of epistemic rationality.

Chapter 1. Setting the Stage

Chapter summary. After having introduced some essential notions and distinctions surrounding the notion of rationality, this chapter presents some important worries and complications in the debate about the normativity of epistemic rationality. Then, I introduce two hypotheses (the Minimal Normative Hypothesis and the Modest Reductionist Hypothesis) that are jointly sufficient for vindicating the normativity of epistemic rationality. That is, if the Minimal Normative Hypothesis and the Modest Reductionist Hypothesis are correct, epistemic rationality is deontically significant. At the end of the chapter, I explain how we can vindicate both hypotheses and I state some basic assumptions that I will make throughout the thesis.

'When I use a word,' Humpty Dumpty said, in rather a scornful tone,
'it means just what I choose it to mean—neither more nor less.'

'The question is,' said Alice, 'whether you can make words mean so many different things.'
'The question is,' said Humpty Dumpty, 'which is to be master—that's all.'
Lewis Carroll, *Through the Looking Glass*

It is controversial whether epistemic rationality is normative. Of course, calling someone "irrational" is not seen as a compliment. Hence, it is commonly assumed that epistemic rationality is desirable, and so it seems that having rational beliefs *qua* rational beliefs is *pro tanto* preferable. But how can we *argue* that epistemic rationality is normative? In this chapter, I explore some worries and complications surrounding this debate. Fortunately, I will come up with a two-steps strategy to vindicate the normativity of epistemic rationality which avoids most of these worries and complications.

A common thread underlying all the parts of this introduction is that we know very little about rationality. Ironically, even the possibility of discovering facts about rationality is subject to controversy. For instance, Errol Lord thinks that we can identify the "real definition of the property of being rational." (Lord 2018, 6) This would entail, among other things, that

the nature of rationality can somehow be discovered. By way of contrast, Alex Worsnip and Michael Titelbaum think that there are different, conventional ways to talk about rationality. When confronted with theories of rationality that conflict with his, Worsnip will often deflate the issue by saying that he is concerned with rationality understood as coherence (Worsnip 2018). In a similar vein, when confronted with objections against some of his views, Titelbaum sometimes says that he can't really be wrong about his account of rationality because that is "just how I [Titelbaum] use the term "rationality"." (Titelbaum Forthcoming, sec. 1) So, the property of rationality would not be something we discover: it is something we agree upon.

Here, we are stuck between Alice and Humpty Dumpty, and it is hard to see who is right and who is wrong. On the one hand, it is tempting to assume that the real property of rationality (or the real definition of the property of being rational) can be discovered or fixed through an essence. But as we will see in section 1.2, there is ample disagreement surrounding the notion of rationality. If rationality has an essence or a real property, this implies that the vast majority of philosophers are conceptually confused when they refer to rationality. This doesn't seem right. On the other hand, if the constraints on the property of rationality are just a matter of convention, it seems that philosophical analysis of this concept will be fairly limited, and that we are very likely to talk past each other. For instance, suppose a philosopher wants to argue that rationality is normative. He or she could simply stipulate that rationality is normative. This would be a cheap demonstration of the view according to which rationality is normative. If the conventionalist is correct, what could be said against such a tactic?

With these problems in mind, I will nevertheless try to make sense of this whole debate and to sketch a convincing strategy for vindicating the normativity of epistemic rationality. Surely, this will be a difficult task, but I hope to make an interesting contribution to this debate.

1.1. Structural and Substantive Requirements of Epistemic Rationality

1.1.1. What Are the Requirements of Rationality?

As I briefly indicated in the above paragraphs, there is little philosophical agreement on the fundamental nature of rationality. But even if rationality had an essence or a nature that we could discover, most philosophers admit that they lack an effective method for determining what are the requirements of epistemic rationality. In other words, regardless of how to settle the *ontological* question of what rationality consists in, there is a persistent *epistemic* problem of how we acquire knowledge of the requirements of rationality. According to John Broome:

> How can we identify requirements of rationality? I wish I could describe a general method of doing so, but I am sorry to say I cannot. I shall defend a number of requirements one by one, on particular grounds that seem appropriate. (Broome 2013, 150)

Broome is here inspired by Thomas Nagel, who suggests that:

> Rationality can be defined only in terms of adherence to rational requirements. One cannot discover or justify the principles which specify those requirements by deriving them from the concept of rationality, since it is precisely those requirements which define the concept, and they must be rendered plausible as requirements independently. (Nagel 1978, 20)

Moreover, it is not always possible to derive requirements of rationality from the necessary conditions on being a rational believer. The mere fact that X is a necessary condition for being a rational believer doesn't entail that X is rationally required. For example, being alive is a necessary condition for being a rational believer, but being alive is not rationally required (and neither is a dead person irrational).[1] Thus, some necessary conditions on being a rational believer have nothing to do with the requirements of rationality. Now, there could be solutions to this problem. For instance, perhaps the facts that X is a necessary condition for being a rational believer and that X is not a necessary condition for being an

[1]See Kiesewetter's (2017, 17–19) discussion of the distinction between the property and the code sense of rationality, and how this affects our capacity to identify the requirements of rationality. See also Broome (2013, chap. 7).

irrational believer together entail that X is rationally required.[2] The point here is that there is an explanatory gap between requirements of rationality and properties of rational agents. We need to say more in order to fill the gap.

Another issue concerns the explanatory role of the requirements of rationality. When philosophers say that X is a requirement of rationality, they are implicitly committed to saying that X is a requirement of rationality with a *distinct* explanatory role—that is, a requirement which explains why some agents are rational and why other agents are irrational. By explanatory role, I am here thinking of an "analytic" explanatory role, in the same sense as the one in which 0+1=1 is part of the analytic explanation of why 2+2=4.[3] Indeed, consider the following examples:

(i) Rationality requires that, if X, then Y.

(ii) Rationality requires that, if X and Z, then Y.

Let's assume that (i) is a requirement of rationality. Since (ii) is entailed by (i),[4] it is natural to conclude that (ii) is also a requirement of rationality. But many philosophers will not accept such a conclusion, since (ii) plays no distinct explanatory role in theories of rationality. That is, since (ii) can be entirely derived from (i), many philosophers will merely accept that (i) is a requirement of rationality, not (ii).

In view of the foregoing, we face important difficulties for identifying the requirements of rationality. First, we lack a consensual method for identifying the requirements of rationality. Second, even if X is a necessary condition for being a rational believer, we can't conclude that X is rationally required. Third, even if we agree that X is a requirement of rationality, and that Y can be derived from X, there will be disagreement on whether Y is a requirement of rationality, because Y will play no distinct explanatory role in the theories of

[2]I thank Daniel Laurier for bringing this possibility to my attention.

[3]For instance, Kiesewetter argues that the following view is unsatisfactory: there are some structural requirements of rationality that are not normative (Kiesewetter 2017, 21). However, he argues that we can explain such requirements in terms of other requirements (e.g., structurally irrational agents do not respond correctly to their reasons) (Kiesewetter 2017, chaps. 9–10). This means that he is implicitly committed to the view that "some requirements of rationality are not normative" is tantamount to "some requirements of rationality *with a distinct explanatory role* are not normative". Hence, he thinks that X is a requirement of rationality insofar as X plays a distinct explanatory role in the theory of rationality.

[4]To be clear: violations of (ii) are also violations of (i), but satisfying (ii) doesn't entail that (i) is satisfied.

rationality. Thus, our philosophical toolbox for identifying the requirements of rationality is fairly limited, to say the least.

1.1.2. Substantive and Structural Requirements

We can remain neutral on the nature of rationality while making a distinction between *types* of putative requirements. Specifically, we can make a distinction between structural and substantive requirements of epistemic rationality.[5] Putative structural requirements of epistemic rationality mostly consist in formal coherence constraints, as in the following:

Consistency. Epistemic rationality requires that, if A believes that P, then it is false that A believes that ~P.[6]

Intra-Level Coherence. Epistemic rationality requires that, if A believes that P_1, believes that P_2, ..., and believes that P_n, then it is false that A believes that $\sim(P_1 \wedge P_2 \ldots \wedge P_n)$.[7]

Inter-Level Coherence (or Reasons Enkrasia). Epistemic rationality requires that, if A believes that he or she has sufficient epistemic reason to believe P, then A believes that P.[8]

By way of contrast, there are putative substantive requirements of epistemic rationality. For instance, it is often suggested that epistemic rationality has to do with responding to sufficient epistemic reasons that the agents have.[9] However, the "responsiveness claim" can be interpreted in at least two different ways, as in the following:

---

[5]It should be noted that some philosophers think that requirements are secondary in the theory of rationality. According to Fogal (m.s.), we should rather focus on sources of normative pressure. For instance, coherence and reasons-responsiveness are two sources of normative pressure. I here leave this possibility aside.

[6]See Broome (2005, 322; 2007a, 355; 2013, sec. 9.2) on Consistency. See also Fogal (m.s.), who thinks we can find counterexamples to all the putative structural requirements discussed in this thesis.

[7]The rational status of Intra-Level Coherence is contentious. Specifically, some solutions to the Lottery Paradox entail that Intra-Level Coherence is not a genuine requirement of epistemic rationality. See notably Demey (2013), Foley (2009) and Sturgeon (2008). I will come back to this problem in chapter 3.

[8]As with Intra-Level Coherence, the rational status of Inter-Level Coherence is also contentious. For example, Coates (2012) and Lasonen-Aarnio (2014; 2015; m.s.) have argued that responding correctly to one's reasons sometimes entail believing "P, but I have sufficient epistemic reason not to believe P", which is an incoherent combination of attitudes. They conclude that such incoherence is not necessarily irrational. See D. Greco (2014), Horowitz (2014a), Kiesewetter (2016), Littlejohn (2015), Titelbaum (2015) and Worsnip (2015) for various responses to this view. I will come back to this problem in chapter 3.

[9]There is also an important debate concerning what it means to *have* a reason. See notably Schroeder (2008; 2011) and Lord (2010; 2018, chaps. 3-4) for various responses to this problem. Also, there is also an important debate concerning what it means to have *sufficient* reason to believe P—see Schroeder (2015) and Lord (2018, chap. 5) for an overview of this debate.

Owned Reasons Thesis. Epistemic rationality requires that, if A has sufficient epistemic reason to believe P, A believes that P. Epistemic reasons are here understood as facts.[10]

Owned Apparent Reasons Thesis. Epistemic rationality requires that, if A has sufficient *apparent* epistemic reason to believe P, A believes that P. Apparent epistemic reasons are here understood as *apparent* facts, which are not necessarily facts.[11]

Typically, reasons for believing P are understood in the narrow sense of ordinary facts or true propositions making it more probable that P (or counting in favour of the conclusion that P).[12] For example, the fact that there is an apple in the room is a reason for me to conclude that there is an apple in the room, provided that I am in a position to perceive the apple. Relatedly, apparent reasons for believing P are apparently true propositions which, if they were true, would count in favour of the conclusion that P. Apparent reasons can be non-factive (i.e., false propositions). For example, if I hallucinate an apple in the room, the content of such a phenomenal state is an apparent reason for me to conclude that there is an apple in the room. Perceiving is to factive reasons as hallucinating is to apparent non-factive reasons.

A quick clarificatory remark: the distinction between reasons and apparent reasons is less stringent than it seems. Indeed, some philosophers understand the notion of reason in a broader sense, which include *facts about appearances*. For example, suppose that I hallucinate that P. My reason to believe that P could be the fact that *it appears to me that P*. With this broader understanding of reasons (e.g., the one that includes facts about appearances), the Owned Reasons Thesis and the Owned Apparent Reasons Thesis are compatible with each other. They just provide different interpretations of the same phenomena.

---

[10]See J. Greco (2005), Goldman (1986) and Littlejohn (2012, esp. 144-55). Kiesewetter (2017, chap. 7) and Lord (2017; 2018) defend a more subtle position: I will come back to such a view in chapter 2. In a capacitist framework, it could be argued that factive reasons amount to factive phenomenal states such as perceptions (by way of contrast with hallucinations)—see Schellenberg (2016).

[11]See Parfit (2011, 34, 111), Conee and Feldman (2004), Feldman (2005), Gibbons (2013), McHugh and Way (2017) and Schroeder (2008; 2011). Such an account could also be understood as a type of substantive internal coherence between (i) *a priori* knowledge and phenomenal experiences and (ii) beliefs or credences—see Turri (2009) or Wedgwood (2017, sec. 0.5).

[12]The claim that reasons for believing P make it more probable that P can notably be found in Lord (2010, 285)., McHugh (2012, 9) or T. Williamson (2002, chap. 9). Some people deny the factivity of reasons and argue that false propositions can be reasons. See Mitova (2017) for an overview of this debate. In this thesis, I assume that reasons are facts, but that false propositions can be apparent reasons. More on this point in chapter 2.

1.2. The "Why-Be-Rational?" Challenges

1.2.1. Which "Why-Be-Rational?" Challenge?

It is now commonly assumed that there are different norms governing beliefs. For example, evaluative norms refer to what bears value, while deontic norms refer to what we are permitted or required to do.[13] For instance, it is commonly assumed that truth is the *evaluative* norm of belief.[14] That is, agents are epistemically *better off* with true beliefs and they are epistemically *worse off* with false beliefs.

However, in some cases, what is true is also very improbable. Tonight's winning lottery combination might be 0-49-21-34-38-11, but such a fact is extremely improbable (as long as they do not announce the results). Hence, there are cases where truly believing P would consist in a leap of faith. It doesn't seem right to believe P based on mere faith. This is why there are *deontic* norms governing belief, that is, norms having to do with permissions and obligations to believe some propositions. These norms allow us to explain why, in some situations, P is true but the agent is not epistemically permitted to believe P. In this thesis, I am interested in the *deontic* normativity of rationality.

Now, there are many interpretations of what the deontic refers to. For instance, some philosophers make a distinction between the deliberative ought and the ought of advice. The deliberative ought has to do with "reasons that matter in first-personal deliberation" (Kiesewetter 2017, 13). By way of contrast, the ought of advice has to do with reasons that matter if I were to receive an advice from a well-informed third party. In this thesis, I am concerned with the deliberative ought and I leave other interpretations of the ought behind.

[13]See McHugh (2012). See also Tappolet (2014) on deontic and evaluative concepts. At this point in the argument, I do not assume that there is a connection between the deontic and the evaluative.

[14]See Ahlstrom-Vij and Dunn (2014) and Goldman (1986; 2015) on veritism. Accuracy-centered epistemology also assumes that truth is the evaluative norm of belief, as in Pettigrew (2016a). See Berker (2013; 2015) for a critique of the above versions of epistemic consequentialism. See Shah (2003), Shah and Velleman (2005) and Whiting (2010) on truth as the aim of belief. See, among others, Gibbard (2012, chap. 4) and Wedgwood (2013b) for discussion concerning true beliefs as correct beliefs. Even Gibbons (2013) and Horwich (2013), who think that epistemic justification (or rationality) is the deontic norm of belief, consider that it would be epistemically irrational for an agent not to care about getting true beliefs. However, there are competing accounts of the evaluative norm of belief—some authors such as Engel (2004) or T. Williamson (2002) have suggested that knowledge is the evaluative norm of belief.

There are at least three formulations of the view according to which epistemic rationality has deontic force, as in the following:

Strict Normativity. If epistemic rationality requires of A to believe P, then A is epistemically required (or permitted) to believe P (Kiesewetter 2017, 20; Lord 2017).

Strict Constitutive Normativity. If epistemic rationality requires of A to believe P, then A is epistemically required (or permitted) to believe P *because* A is rationally required to believe P (Broome 2013, 192).

Contributory Normativity. If epistemic rationality requires of A to believe P, then A has an epistemic reason to believe P (Kiesewetter 2017, 20; Reisner 2011, 36).

Accordingly, there are at least three challenges to the view according to which epistemic rationality is normative (one challenge for each formulation). In this thesis, I try to vindicate Strict Normativity. Plausibly, Contributory Normativity is entailed by Strict Normativity. So, vindicating Strict Normativity probably entails that Contributory Normativity is also correct. However, Strict Normativity is logically weaker than Strict Constitutive Normativity. So, why not try to vindicate Strict Constitutive Normativity instead?

Philosophers like Broome (2013, chap. 11) are interested in Strict Constitutive Normativity because they want to argue that rationality is a *source* of normativity. Specifically, Broome is looking for an inferential, non-derivative vindication of the normativity of rationality. By inferential vindication, I mean that Broome is not satisfied with the plausible pre-theoretical assumption that rationality is normative. Perhaps it is *intuitive* that rationality is normative, but Broome wishes to identify a good *argument* for why rationality is normative (Broome 2013, 204–5). By non-derivative vindication, I mean that Broome is not looking for an explanation in terms of value, optimality, prudence, choiceworthiness, morality, etc. of why rationality is normative. Indeed, even if satisfying the requirements of rationality is value-conducive, optimal, prudent, morally acceptable or the like, Broome will reply that this does not mean that rationality is a source of normativity—rather, this merely indicates that such requirements are normative regardless of whether they are requirements of rationality (Broome 2013, 197–99).

In this thesis, I will not pursue Broome's project. I believe that Broome's distinction between derivative and non-derivative vindication is ambiguous and allows for equivocation. To see this, consider the following explanations of the normativity of rationality:

(i) It is a brute fact or a normative bedrock that rationality is normative.

(ii) Rationality is normative in virtue of explanation X, where explanation X refers to rationality.

(iii) Rationality is normative in virtue of explanation Y, where explanation Y does not refer to rationality.

First, (i) does not provide an inferential vindication of why rationality is normative. Second, (ii) provides a circular or question-begging explanation of why rationality is normative. For instance, suppose that one argues that rationality is normative because rationality consists in responding correctly to the reasons of rationality (by way of contrast with, for instance, the reasons of morality). While such an explanation is inferential and non-derivative, it is either circular or question-begging. Indeed, we still need an explanation of why the so-called reasons of rationality are genuinely normative. Otherwise, we vindicate the normativity of rationality by taking for granted that rationality provides normative reasons!

This leaves us with (iii). However, it seems that (iii) merely provides a derivative explanation of why rationality is normative—for instance, arguing that rational attitudes optimize expected value or are prudent amounts to deriving the normativity of rationality from the normativity of value or prudence. So, (iii) does not seem to provide a non-derivative explanation of why rationality is normative.

Broome argues that there can be satisfactory, non-derivative explanations of why rationality is normative. According to him:

> That there is no explanation that stems from another source does not mean there is no explanation at all.... an account of the nature of rationality might imply that rationality is normative. For instance, it is plausible that rationality is constitutive of agency, so that if we were not rational we would not be agents (Broome 2013, 204).

The problem with the above response is that the distinction between derivative and non-derivative explanations becomes very unclear. Indeed, consider the following putative constitutivist explanation of the normativity of rationality:

Constitutivist Argument. Rational dispositions are constitutive of agency. In other words, there is a constitutive necessary connection between agentivity and rationality. Furthermore, one ought to be an agent. So, one ought to satisfy the requirements of rationality (or at least to be disposed to satisfy the requirements of rationality).[15]

Now, compare the Constitutivist Argument with the following line of reasoning:

Ought to be an Agent Argument. At every possible world, if A is an agent, A has some rational dispositions. So, there is a necessary connection between agentivity and rationality, but such a connection is not constitutive. Furthermore, one ought to be an agent. So, one ought to satisfy the requirements of rationality (or at least to be disposed to satisfy the requirements of rationality).

According to Broome, the Constitutivist Argument provides a non-derivative explanation. However, the Ought to be an Agent Argument provides a derivative explanation. It is unclear why the distinction is normatively significant. First, both arguments rely on the same normative explanans (e.g., one ought to be an agent). If the claim that we ought to be agents is a satisfactory normative explanans in the Constitutivist Argument, why is it an unsatisfactory normative explanans in the Ought to be an Agent Argument? This calls for clarification. Second, the Constitutivist Argument does not necessarily entail that rationality is a *source* of normativity. That is, even if the Constitutivist Argument is correct, the normativity of rationality would be entirely derived from the normative pressure of agency. To put it differently, agents do not have reason to be agents *in addition to* reason to be rational. Rather, they have reason to be agents regardless of whether rationality is constitutive of agency. Hence, talk of non-derivative explanations allows for equivocation: even if the Constitutivist Argument provides a non-derivative explanation of the normativity of rationality, this does not prove that rationality is a *source* of normativity. Again, Broome's understanding of the divide between non-derivative and derivative explanations calls for clarification.

[15]See Southwood (2008).

In order to avoid the above difficulties, I will simply assume that derivative vindications of the normativity of rationality are sufficient. This is why I will try to vindicate Strict Normativity, and not Strict Constitutive Normativity.

1.2.2. Pritchardian Epistemology and the Underdetermination Problem

It could be suggested that the view according to which epistemic rationality is normative is *intuitive*. In accordance with epistemological intuitionism, it could be suggested that the following argument for the normativity of epistemic rationality is correct:

Appeal to Intuitions Argument. If it is intuitive that epistemic rationality is normative, it is highly plausible that epistemic rationality is normative. It is intuitive that epistemic rationality is normative. So, it is highly plausible that epistemic rationality is normative.

H.A. Prichard argued that asking for a principled vindication of theses such as "morality is normative" is improper (H.A. Prichard 1912, 28–29, 33). Plausibly, the same goes for the normativity of rationality. Instead of deriving our obligations from other principles, Prichard suggests that we have a clear and immediate access to such facts through conscious reflection. According to him:

> This apprehension is immediate, in precisely the sense in which a mathematical apprehension is immediate, e.g., the apprehension that this three-sided figure, in virtue of its being three-sided, must have three angle.… we do not come to appreciate an obligation by an argument, i.e. by a process of non-moral thinking, and that, in particular, we do not do so by an argument of which a premise is the ethical but not moral activity of appreciating the goodness either of the act or of a consequence of the act; i.e. that our sense of the rightness of an act is not a conclusion from our appreciation of the goodness either of it or of anything else. (Prichard 1912, 28–29)

Now, Prichard did not provide an explanation of why intuitions are correct in general. According to many philosophers, intuitions are the result of covert processes which give us access to adapted responses to various issues. For instance, it could be argued that intuitions are the result of evolutionary processes. The fact that φ-ing seems intuitively correct to us could be best explained by the fact that our ancestors couldn't have survived without φ-ing. Another possibility is that intuitions are the result of fine-grained psychological mechanisms of information processing. For instance, perhaps our intuitions are the "coarse-grained"

outputs of fine-grained information processing mechanisms in the brain. While we have access to such coarse-grained responses, we usually lack a luminous access to the fine-grained mechanisms underlying such responses.[16]

Be that as it may, I do not find the Appeal to Intuitions Argument convincing. The Appeal to Intuitions Argument is relevant insofar as we share an understanding of the objects we have intuitions about. For instance, we can have relevant intuitions on cooperating with each other because we have a clear picture of what it means to cooperate with each other. This condition is hardly satisfied when it comes to epistemic rationality.

As I indicated in section 1.1, we lack a clear method for identifying the requirements of rationality. This probably explains why there are as many theories of epistemic rationality as there are epistemologists. Todd and Gigerenzer's theory of epistemic rationality has nothing to do with Broome's or with Bayesian accounts of rationality.[17] Many questions surrounding the nature of epistemic rationality are controversial. For instance, does epistemic rationality supervene on an agent's mental states?[18] Assuming that rationality supervenes on the mental, does rationality supervene on an agent's attitudes in a time-slice, or does rationality supervene on an agent's attitudes over time?[19] Should practical and epistemic rationality be theorized independently of each other?[20] Does the context affect the content of the requirements of rationality, or are the requirements of rationality "invariant"?[21] Are non-ideal theories of

[16]See Railton (2014) or T. Williamson (2008, 251–52, 255–57). See also Weinberg et al. (2001) for a critical overview of the methodological role of intuitions in normative philosophy.

[17]Todd and Gigerenzer's (2000) account of epistemic rationality rely on a type of reasons-responsiveness condition for non-ideal agents in terms of fast and frugal heuristics. See also Cherniak (1990). Broome's (2013) account of epistemic rationality rely on structural requirements of epistemic rationality for ideal and non-ideal agents. Objective Bayesianism relies both on structural and substantive requirements of epistemic rationality for ideal agents with great cognitive capacities (see, for instance, Pettigrew (2016a) or J. Williamson (2010))

[18]Broome (2013), Parfit (2011, 34), Reisner (2011), Conee and Feldman (2004) or Wedgwood (2002) think that rationality supervenes on the mental. J. Greco (2005), Goldman (1986) and Littlejohn (2012, 144–45) reject such an assumption.

[19]Hedden (2015a) and Broome (2013) are synchronists about requirements of rationality. Carr (2015), Hlobil (2015) and Podgorski (2016a) are diachronists about requirements of rationality.

[20]Côté-Bouchard (2015; 2016; 2017), Kearns and Star (2008), McGrath and Fantl (2013), Reisner (2009a; 2015) and Rinard argue that the epistemic and the pragmatic should not be theorized independently (we can explain epistemic norms in practical terms, or *vice versa*). See also Levinstein (2017). Ahlstrom-Vij and Dunn (2014), Feldman (2005b), Goldman (2015), Hieronymi (2005), Joyce (1998) and Pettigrew (2016a) think that the epistemic and the pragmatic can (and should) be theorized independently of each other.

[21]Buchak (2013), Clarke (2013), Dorst (2017), Leitgeb (2014a) and McGrath and Fantl (2013) endorse various forms of epistemic contextualism. Pettigrew (2016a) and most veritists endorse epistemic invariantism.

rationality a mere satisfactory approximation of ideal theories of rationality?[22] Which attitudes are governed by requirements of epistemic rationality, credences or beliefs? And what is the relationship between them?[23] What is the relationship between truth and epistemic rationality?[24] And between reasons and rationality?[25] Assuming that rationality has to do with one's reasons, which factors determine what are one's reasons? Is it the facts of my situation or my perspective on the facts of my situation?[26] All these questions (and many others!) are controversial.

To be clear: the Appeal to Intuitions Argument could work if we had a clearer understanding of what epistemic rationality consists in. But this can't be good methodology right now. We would first need to provide a satisfactory account of epistemic rationality. That is, we first need to solve some fundamental issues surrounding the nature of epistemic rationality. Then, perhaps the Appeal to Intuitions Argument could work. But since I firmly believe that we lack a satisfactory account of epistemic rationality, I leave the Appeal to Intuitions Argument aside.

1.2.3. The Objection from Pointlessness

Perhaps we can't appeal to intuitions to vindicate the normativity of epistemic rationality. But perhaps we should not doubt that epistemic rationality is normative in the first place. Why should we doubt that epistemic rationality is normative?

There are at least four important objections against the normativity of epistemic rationality, namely (i) the objection from pointlessness, (ii) the asymmetry/bootstrapping

[22]DiPaolo (m.s.), Staffel (2017) and Todd and Gigerenzer (2000) think that ideal and non-ideal theories can be theorized independently of each other. Smithies (2015) endorse the approximation thesis. See also Talbott (2016, sec. 6.1.A). I will come back to this problem in chapter 2.

[23]Clarke (2013), Dorst (2017) and Foley (2009) think that rational beliefs supervene on rational credences. Easwaran (2015) thinks that rational credences supervene on rational beliefs. Buchak (2014) argues that credences and beliefs should be theorized independently of each other.

[24]Reliabilists like Goldman (1986; 2015) think that substantive requirements of epistemic rationality are truth-conducive. Berker (2013a; 2013b; 2015) rejects epistemic consequentialism. Epistemologists like Feldman (2005b) or Cloos (2015) think that not all epistemic norms are derived from the final value of truth.

[25]Lord (2017) and Kiesewetter (2017) argue that rationality consists in responding to reasons one has. Broome (2013, chap. 5), Kolodny (2007b) and Worsnip (2015; 2016) deny that rationality consists in responding to reasons there are or one has.

[26]See note 9. See Littlejohn (2015) and Titelbaum (2015) on objectivism. See McHugh and Way (2017) on perspectivism.

problem, (iii) the problem of truth-conduciveness and (iv) the objection from unsolvable normative dilemmas. In the remainder of section 1.2, I will present such objections.

Start with the objection from pointlessness. According to this objection, it is unclear whether there is such a thing as epistemic normativity. Let's assume for a moment that epistemic rationality is in the correct relationship with the truth norm of belief (or that we are in a case where satisfying the requirements of epistemic rationality leads to more accurate doxastic states). Even in such a context, some propositions are pointless, unrelated to our interests, plainly uninteresting, and so forth. If P is a true but pointless proposition, not believing P or having a low credence in P does not seem disvaluable, even when having conclusive epistemic reason for P. If this is correct, then having true beliefs *qua* true beliefs is not necessarily *pro tanto* valuable.

To get a better grasp of the problem of pointless truths, here is a toy case. Sarah has conclusive evidence for the following true propositions:

(1) There is an even number of giant sequoias in Muir Woods;
(2) Smoking increases the risk of heart disease.

In determining whether Sarah should believe (1) or (2), most theories of epistemic norms will rely on the likeliness that such propositions are true, or on the weight of the epistemic reasons in favour of such propositions. From an epistemic point of view, practical factors do not play a role in determining whether Sarah should believe those propositions. Specifically, from an epistemic point of view, there is equal epistemic value to have an accurate belief concerning the number of sequoias in Muir Woods or an accurate belief concerning the correlation between smoking and heart disease. This suggests two things, namely that having true beliefs *qua* true beliefs is valuable and that all true beliefs have the same *pro tanto* value (as derived from their epistemic value), as in the following:[27]

Epistemic Value. Epistemic value is a source of *pro tanto* value. Having true beliefs *qua* true beliefs is *pro tanto* valuable and having false beliefs *qua* false beliefs is *pro tanto* disvaluable.

[27]Such a claim has been defended by many authors. See Foley (1993), Goldman (1986), Horwich (2010), Kvanvig (2008; 2009), Lehrer (1990) and Lynch (2004; 2009).

Invariantism. All true beliefs *qua* true beliefs have the same *pro tanto* value and all false beliefs *qua* false beliefs have the same *pro tanto* disvalue.

Now, suppose that Sarah doesn't care about the number of sequoias in Muir Woods and that she has excellent reasons for not caring. Such a belief doesn't satisfy her sense of curiosity, doesn't stimulate her intellect, won't help her reach new interesting conclusions, and so forth. Let's also assume she knows that the same goes for the epistemic community she is a member of. In view of the foregoing, the following seems intuitively correct:

Pointless Truths. From an epistemic or a practical perspective, believing a pointless truth does not seem valuable and disbelieving a pointless truth does not seem disvaluable.

If Pointless Truths is correct, it is not clear why having an inaccurate credence in the proposition "There is an even number of giant sequoias in Muir Woods" should lower her overall utility function.[28] Furthermore, provided that Invariantism is true and that pointless truths have no epistemic or practical value, Epistemic Value is false. Here is why. Assuming that all truths have the same *pro tanto* value, if some attitudes (such as attitudes towards pointless truths) have no *pro tanto* value, this would imply that *all true beliefs have no* pro tanto *value*. Consider the following argument:

(1) All true beliefs *qua* true beliefs have the same *pro tanto* value.

(2) Even from an epistemic perspective, believing a pointless truth has no value (or a value of 0).

(C) Therefore, all true beliefs *qua* true beliefs have no *pro tanto* value (or a value of 0).

Hence, combining Invariantism and Pointless Truths leads to the denial of Epistemic Value. Since most accounts of epistemic rationality have to do with accuracy or truth-conduciveness, we then face an important problem. If Epistemic Value is false, the epistemic norms that are derived from Epistemic Value (such as the consequentialist accounts of epistemic rationality) do not carry weight.

[28]Similar lines of reasoning have been put forth by many authors. See Brady (2009), Côté-Bouchard (2015; 2016), Grimm (2008; 2009), Lockard (2013) and Whiting (2013).

1.2.4. The Objection from Asymmetry and Bootstrapping

Another common objection against the normativity of epistemic rationality is that structural requirements of epistemic rationality lead to bootstrapping. Specifically, not all structural requirements can be wide-scope, and narrow-scope requirements lead to bootstrapping, as in the following:

Asymmetry Problem. Wide-scope structural requirements have the following form: rationality requires that, if A has doxastic attitude α, then A does not have doxastic attitude β. If a requirement of rationality is wide-scope, one can reason from the content of α to dropping β, and *vice versa*. However, not all structural requirements satisfy such a symmetry condition. So, not all structural requirements of rationality are wide-scope.

Bootstrapping Problem. Suppose that some structural requirements are narrow-scope. For instance, if A believes that P, then rationality requires of A not to disbelieve P. Then, A simply needs to form the believe that P to be rationally required not to disbelieve P. But even if an agent believes P, this does not provide him or her a reason against disbelieving P. Perhaps A has no epistemic reason for believing P. In believing P, a reason against disbelieving P appears *out of nowhere*.

Here is why. To begin with, Kolodny introduces the following putative structural requirements of rationality:

I+Wide. Rationality requires that, if A believes that there is conclusive reason to X, A intends to X.

B+Wide. Rationality requires that, if A believes that there is conclusive evidence that P, A believes that P.

I+Wide and B+Wide are inter-level coherence requirements of rationality: they govern the relationship between an agent's first-order attitudes and his or her beliefs concerning what A has sufficient reason to do or believe. The problem with both requirements is that they do not satisfy the symmetry condition. Indeed, consider a case in which A believes that there is conclusive reason to X but lacks the intention to X. As Kolodny indicates:

> One certainly can reason from the content of one's belief that one has conclusive reason to X to an intention to X. So one can rationally resolve the conflict in that way. But one cannot reason from the content of one's not intending to X to a revision of one's belief that one has conclusive reason to X. One cannot reason from the content of one's not intending to X to anything at all, because there is no such content. Not intending to X is simply lacking an attitude. The lack of an

> attitude has no content. (Intending not to X does have content, namely, that one will not X. But intending not to X is not the same as not intending to X) (Kolodny 2005, 527–28)

In *Why Be Rational?*, Kolodny focuses on I+Wide, but the same kind of argument would also hold for B+Wide. Indeed, consider a case in which an agent believes that he or she has conclusive reason to believe P, but lacks the beliefs that P. One can reason from the content of his or her belief that one has conclusive reason to believe P to the belief that P, which resolves the conflict. However, in order to revise one's belief that one has conclusive reason to believe P, one can't reason from the lack of belief that P. So, the symmetry condition is not satisfied.

Kolodny thinks that wide-scope requirements like I+Wide and B+Wide must satisfy the symmetry condition. This leads Kolodny to conclude that there are narrow-scope requirements of epistemic rationality (as in the Asymmetry Problem). To put it differently, if I+ and B+ are requirements of rationality, they must be narrow-scope, as in the following:

I+Narrow. If A believes that there is conclusive reason to X, rationality requires that A intends to X.

B+Narrow. If A believes that there is conclusive evidence that P, rationality requires that A believes that P.

However, there is a problem with I+Narrow and B+Narrow. Such narrow-scope requirements lead to bootstrapping, a process by which requirements appear out of nowhere. Indeed, according to Kolodny:

> Consider the loosely stated rational requirement: 'If you believe that p, then you are rationally required to believe what p entails.' If we read this literally, and assume that the rational requirement has narrow scope, then 'detachment' of the consequent is permitted. From the fact that you believe that p, it follows that you are rationally required to believe what p entails. If we suppose that we have conclusive reasons to comply with rational requirements, then it follows that you have conclusive reason to believe what p entails. And since p entails p, it follows that if you believe that p, then you have conclusive reason to believe that p. Beliefs become 'self-justifying.' (Kolodny 2005, 514)

If a requirement is genuinely normative, such a requirement does not lead to bootstrapping. Accordingly, Kolodny concludes that, since some requirements of rationality lead to bootstrapping, such requirements are not genuinely normative.

Hence, defenders of the normativity of epistemic rationality face a trilemma: either (i) they must explain why structural requirements such as B+ or I+ are not genuine requirements of rationality, or (ii) they must solve the Asymmetry Problem, or (iii) they must solve the Bootstrapping Problem.

1.2.5. The Objection from Truth-Conduciveness

Kolodny (2007b, 230-31) also argues that one doesn't necessarily have an epistemic reason to satisfy structural requirements of epistemic rationality such as Consistency. Kolodny's line of reasoning rely on a specific type of epistemic teleology, namely, veritism. Indeed, he assumes that only true beliefs bear final epistemic value (and only false beliefs bear final epistemic disvalue). Veritists deny the kind of objection that I presented in section 1.2.3, since they assume that all truths, *including pointless ones*, bear final epistemic value.

In a veritist framework, it could be argued that satisfying Consistency does not guarantee a better ratio of true to false beliefs. In fact, some perfectly consistent sets of beliefs are entirely false! Kolodny summarizes his "teleological" argument in the following way:

> From the standpoint of theoretical deliberation—which asks 'What ought I to believe?'—what ultimately matters is simply what is likely to be true, given what there is to go on. (...) [but] formal coherence may as soon lead one away from, as toward, the true and the good. Thus, if someone asks from the deliberative standpoint 'What is there to be said for making my attitudes formally coherent as such?' there seems, on reflection, no satisfactory answer. (Kolodny 2007b, 231)

In other words, if one merely satisfies Consistency, one is not more likely to end up forming true beliefs. So, the mere satisfaction of Consistency does not contribute to reliability, in the sense that it does not improve one's ratio of true to false beliefs. Hence, Kolodny thinks that epistemic teleology is incompatible with the normativity of specific requirements of epistemic rationality such as Consistency.

In the above quote, Kolodny's target are structural requirements of rationality such as Consistency. However, his argument can be extended to the "internalist" account of substantive requirements of epistemic rationality.[29] Recall that, in section 1.1, a distinction has been made between two interpretations of substantive requirements of epistemic rationality. One interpretation is concerned with the factive reasons one has. The other is concerned with the *apparent* reasons one has. Given Kolodny's objection from truth-conduciveness, the latter interpretation is problematic. Here is why.

Defenders of the Owned Apparent Reasons Thesis think that epistemically rational agents ought to respond correctly to apparent (factive or non-factive) epistemic reasons. By way of contrast, defenders of the Owned Reasons Thesis think that epistemically rational agents ought to respond correctly to their factive epistemic reasons. To understand the difference between the Owned Apparent Reasons Thesis and the Owned Reasons Thesis, consider the following two scenarios:

Debby In A Vat. Debby is a brain in a vat. Since reasons are facts, brain-agents lack reasons. Hence, they are ignorant for the most part. However Debby, is different from the others. Debby forms beliefs in accordance with her apparent perceptions, memory and sensations, competently weighs her apparent epistemic reasons, reason correctly and draw clever conclusions concerning her vat-environment. The other brain-agents just form their beliefs randomly. But even if Debby follows such rules, she does not have more true beliefs (or less false beliefs) than the others.

Debby's Doppelganger. Becky is Debby's doppelganger in a counterfactual world. Like Debby, Becky forms beliefs in accordance with her perceptions, memory and sensations, competently weighs her apparent epistemic reasons, reason correctly and draw clever conclusions concerning her environment. Even if Debby and Becky have the same beliefs and credences, most of Becky's beliefs are true, since she is not a deceived brain in a vat.

As we can see from the above cases, if the Owned Reasons Thesis is correct, our two doppelgangers (i.e. agents with exactly the same relevant mental states, experiences, perspectives and epistemic standards) are not necessarily equally rational, since Debby does not have access to factive reasons while Becky's apparent reasons are factive reasons. Hence, it can be rational for Becky to believe P while it is irrational for Debby to believe P. By way of

[29]This possibility is often neglected in the literature. For instance, Kiesewetter (2017) and Lord (2017) do not address this possibility.

contrast, if the Owned Apparent Reasons Thesis is correct, our two doppelgangers are equally rational, since their apparent reasons are identical.

In view of the foregoing, the deontic significance of the Owned Apparent Reasons Thesis is highly problematic in a teleological perspective. According to Kolodny, we ought to satisfy some epistemic requirements because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). But if epistemic permissions and obligations have to do with truth-conduciveness, the Owned Apparent Reasons Thesis doesn't seem to be deontically significant. Indeed, consider the case in which Debby is a brain in a vat who responds to her apparent reasons. Since she only responds to apparent non-factive reasons, her ratio of true to false beliefs is bad. Hence, following a teleological explanation of the deontic significance of epistemic rationality, it seems that Debby's belief-forming process is not normative.

In summary, Kolodny's objection from truth-conduciveness is problematic for (i) structural requirements of epistemic rationality such as Consistency and (ii) substantive "internalist" requirements of epistemic rationality such as the ones underlying the Owned Apparent Reasons Thesis.

1.2.6. The Objection from Unsolvable Normative Dilemmas

The last objection against the normativity of rationality stems from the following plausible metaethical principle:

No Conflict. Deontically significant requirements do not conflict with each other, in the sense that they do not lead to unsolvable normative dilemmas.[30]

Here is why No Conflict is problematic when it comes to vindicating the normativity of epistemic rationality. Several authors such as Maria Lasonen-Aarnio and Allen Coates have recently suggested that, in cases where one has misleading higher-order evidence, it can be rational for one to believe that one's epistemic reasons (including deductive reasoning and evidence) do not support believing P. However, recall that there are rational false beliefs. So, perhaps one is nonetheless rational in believing P. In such a case, one could be rational in

[30]Chang (2001), Bélanger (2011) support No Conflict. Christensen (2010, 212), Sinnott-Armstrong (1996) and Williams (1965) argue that No Conflict is implausible. No Conflict can also be interpreted as the following plausible principle of deontic logic: ~(O(φ)^O(~φ)).

believing P *and* respond correctly to his higher-order evidence in concluding that his epistemic reasons do not support that conclusion. This means that structural requirements of epistemic rationality such as Inter-Level Coherence can conflict with substantive requirements such as responding correctly to reasons one has.

Lasonen-Aarnio and Coates solve this conflict by arguing that Inter-Level Coherence is not a genuine requirement of epistemic rationality. (Coates 2012, 113–15; Lasonen-Aarnio 2014; Lasonen-Aarnio m.s.). Alex Worsnip denies such a conclusion and argues that evidence-responsiveness and inter-level coherence "are, properly understood, fundamentally different kinds of normative claims, such that they should not be stated using the same normative concept" (Worsnip 2015, 6). As I will explain in chapter 3, none of the above positions is comfortable or sits well with pre-theoretical assumptions concerning epistemic rationality. Furthermore, the only other alternative is to argue that there are cases where the demands of structural rationality and substantive rationality cannot be simultaneously satisfied. But this amounts to denying No Conflict. Once again, we face an uncomfortable option, which is to deny plausible metaethical principles concerning permissions and obligations.

1.3. Existing (Unsatisfactory) Responses to the Above Challenges

There are many known arguments in favour of the normativity of epistemic rationality. Some authors have argued that the epistemic normativity of rationality has to do with agency (Southwood 2008). Others have argued that particular requirements of epistemic rationality such as Consistency can be derived from the correct response to a priori falsehoods (Reisner 2011). Others have offered new accounts of the scope of structural requirements of epistemic rationality in order to avoid the asymmetry and boostrapping objections.[31] While these arguments can solve some of the above objections, they are unable to solve all of them.

One recent strategy comes from Benjamin Kiesewetter and Errol Lord, who argue for the following views:

[31]See notably Broome (2007b), Brunero (2010), Way (2009; 2010a; 2010b; 2011; 2012), Shpall (2013) and Titelbaum (2015) on such strategies.

Attending Condition. Epistemic rationality has to do with the propositions agents *attend to*, which excludes pointless propositions. That is, rational agents are epistemically permitted to ignore pointless propositions.[32]

Robust Reductionism. Putative structural requirements of epistemic rationality can be explained in terms of failures to satisfy substantive requirements of epistemic rationality. For instance, if one violates putative requirements such as Consistency or Inter-Level Coherence, one did not respond correctly to one's reasons (or one's apparent reasons).[33]

Owned Reasons Thesis. Epistemic rationality requires that, if A has sufficient epistemic reason to believe P, A believes that P.[34]

The Attending Condition answers the problem of pointlessness by limiting epistemic normativity to significant propositions. Robust Reductionism answers Kolodny's objections from bootstrapping and truth-conduciveness, since it entails that structural requirements play no distinct explanatory role in theories of epistemic rationality. That is, we can remain neutral on whether structural requirements of epistemic rationality are normative and determine whether substantive requirements are normative, since structural requirements play no distinct explanatory role. For similar reasons, Robust Reductionism answers the objection from unsolvable normative dilemmas: there can't be a conflict between structural and substantive requirements, because agents who satisfy substantive requirements automatically satisfy structural requirements. Finally, the Owned Reasons Thesis provides a positive answer to the Why-Be-Rational? challenge. Under the assumption that the Owned Reasons Thesis is correct, we have a fairly straightforward explanation of why an important component of rationality is normative: since reasons are the canonical normative unit, agents necessarily have normative reasons to be rational.

Overall, I think that Kiesewetter and Lord are in the right direction. Still, I am unsatisfied with many aspects of their arguments. First, Robust Reductionism relies on the assumptions that epistemic reasons are impermissive (in the sense that they never warrant more than one doxastic attitude towards P)[35] and that higher-order reasons are always coherent

[32]See Kiesewetter (2017, 182–84) and Harman (1999).
[33]See Kiesewetter (2017, 180–85 and chap. 9) and Lord (2017, 15).
[34]See Kiesewetter (2012; 2017) and Lord (2010; 2014; 2017). Reasons can here be understood in the narrow sense of ordinary facts counting in favour of the conclusion that P, but can also be understood in a broader sense which includes facts about appearances. I come back to this point in chapter 2, and I explain why this second interpretation of epistemic reasons raises a worry in view of Kolodny's objection from truth-conduciveness.
[35]See notably Kiesewetter (2017, 180–85). More on this point in chapter 4.

with first-order reasons.[36] These assumptions are highly contentious. If we are to offer a satisfactory argument in favour of Robust Reductionism, we need to take seriously the possibilities that epistemic reasons are permissive and that higher-order reasons can be incoherent with first-order reasons.

Second, the Owned Reasons Thesis comes at an important cost. Recall that, typically, theories of epistemic rationality (such as the Owned Apparent Reasons Thesis) have to do with responding correctly to apparent reasons agents have. Kiesewetter's and Lord's theories of epistemic rationality have to do with responding correctly to reasons agents have. Reasons are a canonical normative unit, but it is far from clear that apparent reasons are also normative. A satisfactory vindication of the normativity of epistemic rationality should not preclude the Owned Apparent Reasons Thesis.

1.4. An Itinerary to the Normativity of Epistemic Rationality

1.4.1. The Normative and the Reductionist Hypotheses

As I have just indicated, I believe that Kiesewetter and Lord have identified a promising strategy for vindicating the normativity of epistemic rationality, but many of their assumptions are problematic. First, a good vindication of the normativity of epistemic rationality should not preclude the normativity of the Owned Apparent Reasons Thesis. Indeed, in accordance with the Owned Apparent Reasons Thesis, many philosophers think that rationality has to do with internal states such as *apparent* reasons one has.

Kiesewetter and Lord claim that their theories of epistemic rationality make room for apparent reasons. Indeed, they argue that facts about appearances can be part of one's reasons. They are a special type of reasons, namely, reasons that consists in facts about mental states. Accordingly, they think that we do not need to vindicate the normativity of apparent reasons. All we need to do is vindicate the normativity of reasons, which also vindicates the normativity of reasons that consists in facts about appearances.

[36]See Lord (2017, 15) and Kiesewetter (2017, 250–54). See Worsnip (2015) for discussion. I will come back to this assumption in chapter 3.

In chapter 2, I will explain why their strategy is problematic. In a nutshell, here is why. If we take Kolodny's objection from truth-conduciveness seriously, it is unclear why facts about appearances are normative. Indeed, responding to such facts does not guarantee that an agent will end up with a good ratio of true to false beliefs. So, in order to provide a satisfactory answer to Kolodny's objection, we still need to explain why appearances (or facts about appearances) are normative. In other words, even if we think that Kiesewetter and Lord's account of reasons is correct, Kolodny's objection from truth-conduciveness should still worry us. In a teleological perspective, reasons that consist in facts about appearances face the same kind of challenge as apparent reasons: neither seem particularly truth-conducive.

I'll call the view according to which both reasons and apparent reasons are normative the Minimal Normative Hypothesis:

Minimal Normative Hypothesis. Agents ought to respond correctly to (apparent) sufficient epistemic reasons they have. Responding correctly to (apparent) reasons one has is deontically significant.

I will argue for such a view in chapter 2.

Second, we need to take a closer look at epistemic permissiveness (the view according to which epistemic reasons are permissive) and epistemic incoherentism (the view according to which higher-order reasons can be incoherent with first-order reasons). Epistemic permissiveness and epistemic incoherentism are important obstacles to vindicating Robust Reductionism. After a careful examination of epistemic incoherentism and epistemic permissiveness in chapters 3 to 5, I will come to the conclusion that a modest version of reductionism is correct, as in the following:

Modest Reductionist Hypothesis. In ideal theories of epistemic rationality, putative structural requirements such as Inter-Level Coherence, Intra-Level Coherence or Consistency have no distinct explanatory role when compared with substantive requirements of epistemic rationality.

1.4.2. The Game Plan, Chapter by Chapter

Chapter 2 is devoted to vindicating the Minimal Normative Hypothesis. As I explained in section 1.2.5, the main problem with the Owned Apparent Reasons Thesis is that responding correctly to apparent reasons one has doesn't guarantee that one will end up with a

good ratio of true to false beliefs. We lack a goal-oriented explanation of why one ought to respond correctly to apparent reasons one has. The challenge, then, is to provide a vindication of the deontic significance of the Owned Apparent Reasons Thesis *in a teleological perspective*. By way of contrast, responding correctly to reasons one has is generally truth-conducive. So, defenders of the Owned Reasons Thesis still seem in a better position to vindicate the normativity of epistemic rationality. So, how do we solve this challenge?

My solution to this challenge relies on a new theory of epistemic norms, the epistemic theory of the second best, which aims at explaining what agents should do when the best option is unavailable to them. In such cases, agents should opt for the best option *available*, or the *second-best* option. At the end of chapter 2, I will explain how the deontic significance of epistemic reasons and the epistemic theory of the second best support the Minimal Normative Hypothesis.

Chapter 3 supports the Modest Reductionist Hypothesis by arguing that Inter-Level Coherence and Intra-Level Coherence play no distinct explanatory role in theories of epistemic rationality. This conclusion will emerge from a close examination of the debate surrounding the rational status of Inter-Level Coherence and the role of reasons in theories of epistemic rationality.

Still, perhaps Consistency plays a distinct explanatory role in theories of epistemic rationality. We lack an explanation of why inconsistent agents necessarily failed to respond correctly to their epistemic reasons. Enters chapter 4. In this chapter, I examine the relationship between Consistency and Extreme Reasons Permissiveness, which states that there are cases where both believing P and disbelieving P are warranted by a body of epistemic reasons. In this chapter, I argue for a conditional: if Extreme Reasons Permissiveness is false, then Consistency plays no distinct explanatory role in theories of epistemic rationality. Naturally, many philosophers have good arguments against Extreme Reasons Permissiveness. In response to their objections, I introduce the Permissive Epistemic Standards Thesis, the view which roughly states that an epistemically rational agent's epistemic reasons are subjectively mediated through some rational epistemic standards, and

that there are incompatible but equally rational epistemic standards available to agents. I argue that we have greatly underestimated the possibility that Consistency plays a distinct explanatory role in the combinations of rational epistemic standards one can entertain.

How can we argue that the Permissive Epistemic Standards Thesis is false? Answering this question is the topic of chapter 5. In this chapter, I will argue that the Permissive Epistemic Standards Thesis is false when it comes to ideal theories of epistemic rationality. I will roughly argue that the view according to which there are equally reliable but incompatible epistemic standards entail the existence of a unique and more reliable epistemic standard. My strategy relies on a well-known theorem in the social choice literature, namely, Condorcet's Jury Theorem. This gives rise to an important problem for those who argue that epistemic standards are permissive, since the reliability criterion is incompatible with such a type of Permissiveness. At the end of the chapter, I discuss the limits of my argument and I respond to various objections against it.

Taken together, chapters 3 to 5 confirm the Modest Reductionist Hypothesis. This completes the argument.

1.5. Some Background Assumptions

There is no free lunch in philosophy and this thesis is no exception. Throughout the thesis, I will make a number of important assumptions for which I can't provide a principled vindication. In this section, I here present the most important assumptions I will make and where they will be mobilized in my arguments.

1.5.1. The Practical and the Epistemic

An important debate that I will not address in this thesis concerns the relationship between the practical and the epistemic. With respect to this debate, I will make two important assumptions. First, without proof to the contrary, I will take for granted that epistemic and practical requirements of rationality are not radically distinct. For example, *ceteris paribus*, if expected utility maximization is correct in the practical realm, then expected epistemic utility

maximization should also be correct in the epistemic realm. This assumption is tacit in my discussion of the epistemic theory of the second best (as in chapters 2 and 5).

Second, I will ignore a possible view concerning epistemic normativity. There are many competing accounts of epistemic normativity. The most popular views are: (i) the epistemic value of having true/justified beliefs necessarily contributes to final value, and such a source of final value is distinct from practical value, (ii) the epistemic value of having true/justified beliefs necessarily contributes to final value, but the notion of "epistemic value" is a *figure of speech*, in the sense that it refers to the practical value of having some cognitive goods, (iii) the existence of pointless truths entails that not all true/justified beliefs bear final value.

If (iii) is correct, this means that the normativity of epistemic rationality is merely a datum and does not apply to pointless propositions. I here ignore (iii), and this is a limit of my argument. In any case, this is a much more general problem which goes well beyond the scope of my thesis.

1.5.2. Substantive Rationality and Responsiveness

I will assume that epistemic rationality has something to do with substantive requirements such as responding to (apparent) reasons one has. Specifically, I deny that epistemic rationality merely has to do with structural constraints such as Intra-Level Coherence, Inter-Level Coherence or Consistency. This assumption is central to the argument of my thesis and will play a role in every chapter.

First, a clarificatory remark: as I indicated in the first pages of this chapter, the distinction between structural and substantive accounts of epistemic rationality might be merely terminological. For instance, defenders of the view according to which epistemic rationality merely has to do with structural requirements could very well acknowledge that there is a more substantive sense of what rationality consists in, but prefer to use the term in a narrower (or minimal) sense. As Worsnip indicates:[37]

[37]See Scanlon (1998, 23–30) on a similar point.

> [Several structuralists] are happy to acknowledge that there's a potentially legitimate usage of 'requirement of rationality' whereby there are requirements of rationality to respond appropriately to reasons; still, they don't want to lose sight of the fundamental distinction Broome is after, and so they use 'requirement of rationality' in a narrower way to refer only to requirements pertaining to rationally (im)permissible combinations of attitudes. (Worsnip 2018, 63)

If the distinction between structural and substantive rationality is merely terminological, then all I can say is that I am interested in the substantive usage of epistemic rationality. Specifically, I am interested in ordinary types of epistemic irrationality that a structuralist can't explain. For instance, consider the following case:

Planet of the Ricks. Rick believes that, if the theory of evolution is true, then some great grandfathers of human children were apes. However, he claims that no great grandfather of human children was an ape. This leads him to conclude that the theory of evolution is false. Rick's reasoning is perfectly *valid*. His beliefs are perfectly *consistent*. Even better: his beliefs satisfy Inter-Level Coherence and Intra-Level Coherence! Still, Rick has numerous salient perceptions and memories indicating that his reasoning is based on wildly implausible premises. For example, Rick has heard a very large number of arguments against his premise that, if the theory of evolution is true, then some great grandfathers of human children were apes. While he remembers most of these arguments, *Rick simply ignores them*.

Let's ignore all the other attitudes Rick might have and evaluate his rationality with respect to his beliefs, perceptions and memories mentioned in the above case. In a structuralist framework, then, Rick appears to be epistemically rational. After all, his beliefs are jointly consistent, he satisfies all formal coherence requirements of rationality (such as Intra-Level Coherence or Inter-Level Coherence), and responding to reasons such as salient perceptions is not rationally required. However, I take it as a datum that Rick is a textbook case of irrationality in the ordinary (or broad) sense. He is a paragon of dogmatism or delusion. If this is correct, Rick's irrationality has to be explained in terms of his failure to respond to his reasons. This is why I am inclined to endorse the broader usage of the concept of epistemic rationality, which has something to do with responding to reasons one has.

The structuralist could then reply that Rick violates Inter-Level Coherence. Indeed, perhaps Rick believes that the arguments he has heard are conclusive. In such a case, Rick

would believe that there is a conclusive epistemic reason against some of his first-order beliefs, which violates Inter-Level Coherence.[38] However, we need not make such an assumption. Perhaps Rick has heard a very large number of arguments against his premise without forming a higher-order judgment concerning the conclusiveness of such arguments. Rick can't violate Inter-Level Coherence if he doesn't entertain higher-order beliefs concerning the conclusiveness of his epistemic reasons.

Now, the problem is that some philosophers do not think that the distinction between structural and substantive epistemic rationality is merely terminological. A well-known defense of the severe separation between reasons and rationality can be found in Broome's *Rationality Through Reasonin*g (Broome 2013, Chapter 5). Broome denies that responding correctly to reasons one has is a necessary or a sufficient condition for being rational.[39] Broome first presents his "quick objection": while reasons agents have might require them to believe P, if they *ignore* that their reasons require them to believe P, it is not irrational for them not to believe P (Broome 2013, 74–78). In such a case, Broome thinks that agents are not rationally required to believe P, since ignorance of reasons is a rational excuse for not believing P. This means that responding correctly to reasons is not required for being rational.

Broome then considers the possibility that rationality has to do with responding to a specific class of reasons—namely, *attitudinal* reasons. By attitudinal reasons, Broome means "reasons that consist in attitudes" (Broome 2013, 75). After all, considering such a possibility is in accordance with Broome's view that rationality supervenes on mental states (Broome 2013, 89, 151). However, Broome thinks that, in some situations, responding to attitudinal reasons leads to bootstrapping. Even if an agent believes P and (~P v Q), this does not give him or her a reason to believe Q. For instance, an agent could have no epistemic reason in favour of his or her beliefs that P and (~P v Q). In believing P and (~P v Q), a reason in favour of Q would appear *out of nowhere*. Broome takes such a bootstrapping result to be nonsensical

[38]Alternatively, perhaps Rick believes that the arguments he has heard are inconclusive. In such a case, one could argue that his belief "screens the epistemic reasons," in the sense that such a higher-order belief defeats or undermines the arguments he has heard.

[39]As noted by Kiesewetter (2017, 161–62), Broome sometimes conflate (i) responding to reasons there are with (ii) responding to reasons one has. I here assume that Broome means "responding to reasons one has".

(Broome 2013, 81–82). This leads him to deny that (i) rationality is identical to responding correctly to attitudinal reasons one has and (ii) responding correctly to attitudinal reasons one has is sufficient for being rational.[40]

I'll make two brief remarks on Broome's argument against the connection between reasons and rationality. First, following Kiesewetter, Broome's quick objection could be avoided by making a distinction between available and unavailable reasons.[41] Indeed, substantive theories of epistemic rationality claim that rational agents respond to reasons available to them. While one could be rational in ignoring unavailable reasons, it is far less clear that one is rational in ignoring available reasons. Suppose that Rick has no evidence that his perception is flawed and that his perceiving a dog in front of him is maximally salient (in the sense that he is currently unable to think about anything else). Plausibly, Rick's belief that there is no dog in front of him would be irrational, notably because the reasons against such a conclusion are salient and available to him.

The problem is that the notion of ignorance is unclear in Broome's argument. Ignorance can refer to the lack of access to information, but it can also refer to something intentional such as the refusal to take notice of information. If the reasons are outside one's ken (in the sense that one lacks access to the reasons), one is not required to respond to such reasons because one *doesn't have* such reasons. But if the reasons are within one's ken, ignoring the reasons (in the sense that one refuses to take notice of the information) is not rationally excusable. For instance, it is not rationally excusable for Rick to ignore the reasons that are available to him.

Second, in evaluating whether rationality consists in responding to attitudinal reasons, Broome considers only one class of attitudes—namely, beliefs. The problem is that there are other attitudes that could count as reasons. Phenomenal experiences (such as perceptions, memories, seemings, insights, emotions and the like) could be legitimate attitudinal reasons.

[40]A quick clarificatory remark: it is still unclear whether Broome's argument is compatible with the claim that if an agent is rational, then he or she has responded correctly to his or her attitudinal reasons (this is what he calls "Limited Entailment"). On page 79, he claims that his objection from bootstrapping does not affect Limited Entailment, but on page 82 he claims to have shown that Limited Entailment is empty.
[41]See note 39.

Broome does not rule out the view that rationality has to do with responding to phenomenal experiences. This means that, even if Broome is right about responding correctly to beliefs, rationality could nevertheless have to do with responding to other attitudinal reasons.

In summary, structuralists cannot explain why dogmatism and delusion are epistemically irrational. Since I am interested in a broad notion of epistemic rationality, which should be able to explain why dogmatism and delusion are irrational, I do not endorse structuralism about epistemic rationality. Furthermore, Broome's argument for the separation between reasons and rationality is inconclusive. His claim that ignorance of reasons provides rational excuses does not entail that reasons and rationality belong to separate normative domains.

1.5.3. Representing the Weight of Reasons

I will assume that the weight of epistemic reasons is represented by *epistemic probabilities*, understood as the probabilities warranted by an agent's body of epistemic reasons (or apparent reasons). Accordingly, I will also assume that fallible reasons to believe P warrant an epistemic probability of less than 1 in P, and infallible reasons to believe P warrant an epistemic probability of 1 in P. Also, while rational credences might not be *identical* to epistemic probabilities, they *track* epistemic probabilities.[42]

The exact interpretation of what epistemic probabilities are is left open. Some philosophers think that epistemic probabilities are objective or mind-independent (such as frequencies, degrees of causal connection, chances and the like). Others think that epistemic probabilities are internal states (such as rational credences). I will not endorse a specific interpretation of what epistemic probabilities are. However, one thing that I wish to stress is this: I do not believe that, in order to have reasons, agents need to entertain fine-grained credences. So, perhaps epistemic probabilities are internal states, but this doesn't entail that rational agents who have reasons to believe P necessarily entertain a credence in P. The "internal" epistemic probability might be implicit.

[42]In the case where epistemic reasons are permissive (see chapters 4 and 5), we can assume that rational credences track some of the epistemic probabilities warranted by a body of epistemic reasons.

This representation of the weight of epistemic reasons is less than ideal. The probabilistic representation of reasons raises methodological difficulties. It is not always clear how we should represent perceptual learning, defeaters and undermining evidence in a probabilistic framework (Christensen 1992; Pryor 2013; Weisberg 2015). While we should take these difficulties seriously, there are two reasons why I maintain a probabilistic representation of the weight of epistemic reasons in my thesis.

First, as we will see in chapter 3, some authors arguing against Inter-Level Coherence as a requirement of epistemic rationality make use of the probabilistic representation of fallible reasons. Since my goal is to address other arguments found in the literature, it seems justified to make use of the probabilistic representation of reasons. Second, even if the probabilistic representation of reasons is limited and problematic, understanding the type of results we can get in this framework could eventually help us to develop similar arguments in other frameworks. So, even if this is not the most adequate representation of the weight of reasons, it is worth considering what results we reach through such a representation.

1.5.4. Counterfactual and Actual Reliability

The last assumption that I will make concerns the relationship between epistemic rationality and reliability. In this thesis, I will assume that the following is correct:

Reliability Criterion. In the right conditions, if A is ideally rational, A satisfies some available epistemic standards that optimize his or her ratio of true to false beliefs (and such standards lead A to reach the right answer more than 50% of the time).

This assumption is logically weaker than orthodox reliabilism, according to which reliability is necessary and sufficient for epistemic rationality (Goldman 1986). It is also compatible with various forms of internalism concerning epistemic rationality. Indeed, consider the case in which Debby is a brain in a vat who responds to her apparent reasons. In such a case, Debby is *actually* unreliable but *counterfactually* reliable. That is, *if she found herself in the right conditions*, her ratio of true to false beliefs would be good. However, since she only responds to apparent non-factive reasons, her ratio of true to false beliefs is bad. Hence, following a teleological explanation of the deontic significance of epistemic rationality, it seems that

Debby's belief-forming process is not normative. Of course, in normal conditions (for example, if Debby were taken out of the vat and followed the same belief-forming processes), she would be reliable. So, it seems that a teleological explanation of the deontic significance of rationality yields the following verdict: responding to apparent reasons is merely *counterfactually* normative, and responding to factive reasons is genuinely normative.

As we can see, the assumption that epistemically rational agents are counterfactually or actually reliable does not solve Kolodny's objection from truth-conduciveness. Indeed, Kolodny can gladly concede that requirements of epistemic rationality secure counterfactual reliability, but argue that being counterfactually reliable *is not good enough*. Still, as I will explain in chapter 2, this assumption is important for solving Kolodny's challenge from truth-conduciveness. It also plays a role in my argument against the Permissive Epistemic Standards Thesis, presented in chapter 5.

On to the argument.

Chapter 2. Apparent Reasons Are Deontically Significant

Chapter summary. The deontic significance of epistemic reasons (understood as true propositions which count in favour of the conclusion that P) is fairly uncontroversial. After all, reasons are the canonical normative unit. But what about apparent reasons? This chapter argues that apparent epistemic reasons to believe P (understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P) are deontically significant. That is, an agent's epistemic permissions and obligations are affected by the balance of apparent epistemic reasons he or she has. The chapter also responds to two objections against the normativity of apparent reasons. According to the first objection, epistemic norms bear deontic significance because they are a sufficiently good approximation of epistemic ideals. However, responding to apparent reasons (and, especially, to non-factive apparent reasons) is not part of ideal epistemic scenarios. According to the second objection, epistemic norms bear deontic significance because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). However, responding to apparent non-factive reasons is not truth-conducive. These two objections entail that responding to apparent non-factive reasons does not bear deontic significance. The normativity of apparent reasons and both objections are considered in a new framework (i.e. second-best epistemology), which takes into account the legitimate constraints on available belief-forming processes.

Becky responds correctly to her epistemic reasons—for instance, she forms beliefs (or credences) in accordance with her perceptions, memory and sensations. By way of contrast with agents who form their beliefs randomly, her ratio of true to false beliefs is very good (and if Becky entertained credences instead of beliefs, they would be highly accurate). She finds herself in what we can call the "Good Case." Most philosophers think that there is something epistemically commendable with Becky's belief-forming processes. In fact, many philosophers accept something like the following:

Datum. Epistemic reasons, understood in the narrow sense of *ordinary facts* counting in favour of the conclusion that P,[43] have deontic force. That is, an agent's epistemic permissions

[43]See Sylvan (2016a; 2016b) on epistemic reasons. Some people deny the factivity of reasons and argue that false propositions can be reasons. See Mitova (2017) for an overview of this debate. I here assume that reasons in the narrow sense are true propositions, but that false propositions can be apparent reasons. However, my argument could be reformulated to square with other accounts of reasons. For instance, suppose false propositions can be reasons. Then, my argument in this chapter could be reformulated as follows: factive reasons (e.g. reasons that consist in facts) are deontically significant. If factive reasons are deontically significant, then so are non-factive

and obligations are affected by the balance of epistemic reasons he or she has.[44] The deontic significance of epistemic reasons can be restricted to a specific understanding of permissions and obligations (e.g. the deliberative ought, the ought of advice, etc.).[45]

The Datum explains why there is something commendable with Becky's belief-forming processes. I here assume that the Datum is correct. It has a large following and I will not defend it here. Those who reject it can conditionalize my argument on the assumption that the Datum is correct.[46]

The Datum merely says that reasons have deontic force.[47] It doesn't say that *only* reasons are deontically significant. For instance, some philosophers have argued that some false propositions also have deontic significance. To see why, consider the following case:

Bad Case. Debby is Becky's doppelganger, but she is a deceived brain in a vat. Since reasons in the narrow sense are ordinary facts, Debby lacks reasons. However, Debby forms beliefs in accordance with her non-factive experiences. In short, her beliefs are sensitive to *appearances*.

---

reasons. So, non-factive reasons also bear deontic significance.

[44]I am glossing over some subtleties here. Agents might be required to respond correctly to their epistemic reasons insofar as other conditions are fulfilled, such as caring about P, explicitly wondering whether P, considering that P is not a pointless proposition, and the like. For the sake of simplicity, I will assume in the remainder of this chapter that such conditions are always fulfilled. There is also an important debate concerning what it means to *have* a reason. This is an orthogonal issue that I do not wish to address here. See notably Schroeder (2008; 2011) and Lord (2010) for various responses to this problem. Finally, there is also an important debate concerning what it means to have *sufficient* reason to believe P, and whether having sufficient reason to believe P makes it merely *permissible* or *required* for agents to believe P. Again, I do not wish to address this issue here—see Kiesewetter (2017, sec. 7.7), Kelly (2014), Kroedel (2011), Schroeder (2015), Sylvan (2015a) and White (2014) on these debates. [Redacted sentence]

[45]The deliberative ought has to do with "reasons that matter in first-personal deliberation" (Kiesewetter 2017, 13). By way of contrast, the ought of advice has to do with reasons that matter if I were to receive an advice from a well-informed third party. See also Wedgwood (2015) on the various types of ought. These types of ought are not necessarily incompatible with each other. Some authors are interested with the natural process by which agents come to acquire knowledge, while others are concerned with how agents should deliberate. This gives rise to various understandings of "ought" that are perfectly legitimate depending on the kind of problem we are interested in. See Littlejohn (2012), Sylvan (m.s.) or Worsnip (2015; 2016), for example.

[46]Some people might offer a debunking argument of why Becky's belief-forming process seems normative. For instance, it could be argued that A ought to believe P if and only if P, or that A ought to believe P only if A is in a position to know P. Accordingly, reasons do not necessarily bear deontic significance, since Becky can respond correctly to her reasons and form reasonable false beliefs (or fail to form true beliefs). People who endorse such a line of reasoning can then offer the following debunking argument: it appears to us that reasons-responsiveness bears deontic significance because we conflate obligations and excuses. See Littlejohn (2012) or Williamson (forthcoming) on this line of reasoning.

[47]Hereafter, when talking about reasons and apparent reasons in this chapter, I drop the adjective "epistemic."

One interpretation of the Bad Case is that, while Debby lacks reasons, she has *apparent* reasons to believe various propositions. Apparent reasons are here understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P.[48] An alternative interpretation is that, while Debby lacks reasons in the *narrow* sense, she has reasons in a *broad* sense. Facts about appearances (e.g., the fact that it appears to Debby that P) could be her reasons to believe various propositions. So, Debby could have reasons, but not in the narrow sense of ordinary facts.[49] Whether Debby has apparent reasons or reasons that consist in facts about appearances is a separate issue that I do not wish to address here. For simplicity, I will assume that Debby has apparent reasons to believe various propositions and leave aside the second possibility.

In accordance with such an interpretation of the Bad Case, some philosophers endorse the following view:

Deontic Significance of Apparent Reasons. Apparent reasons (understood as apparently true propositions which, if they were true, would count in favour of the conclusion that P) have deontic force. That is, an agent's epistemic permissions and obligations are also affected by the balance of apparent reasons he or she has.[50]

Naturally, when apparently true propositions are true, apparent reasons are also reasons. So, following the Datum, apparent reasons are deontically significant when appearances are factive. However, in the Bad Case, the factivity condition is not satisfied. Thus, it is unclear whether non-factive apparent reasons have deontic force.

The Deontic Significance of Apparent Reasons matters for vindicating the Minimal Normative Hypothesis (see section 1.4.1). Indeed, while reasons in the ordinary sense are a canonical normative unit, it is far from clear that apparent reasons are also deontically

[48]There is an important debate concerning the nature of apparent reasons. See Sylvan (2015b) for an account of apparent reasons in terms of competence and attraction. See Parfit (2011, 34) for an account of apparent reasons in terms of beliefs. With respect to this debate, the counterfactual account of apparent reasons I endorse is not entirely noncommittal. In any case, the debate on the nature of apparent reasons is tangential to the project of this chapter. What matters here is that a massively deceived brain in a vat like Debby has apparent reasons to believe various propositions. Everyone will agree on that, regardless of the specific account of apparent reasons they endorse.

[49]See notably Kiesewetter (2017, 171–74). See also Lord (2017).

[50]See, among others, Lord (2017), Kiesewetter (2017, 171–74) and Parfit (2011, 34, 111). To be specific: Lord and Kiesewetter endorse the deontic significance of reasons that consist in facts about appearances. But as I indicated in the previous paragraph, I leave aside this interpretation of the Bad Case.

significant. Those who think that rationality has to do with reasons in the narrow sense can provide a good explanation of why rationality is normative: reasons-responsive agents necessarily have normative reasons (Kiesewetter 2017, chap. 7; Lord 2017; 2018). However, those who think that rationality has to do with reasons in the broad sense, or with apparent reasons, face a challenge: they need to explain why apparent reasons are normative.[51] Hence, the Deontic Significance of Apparent Reasons is the main obstacle to vindicating the Minimal Normative Hypothesis.

This chapter offers an explanation of why apparent reasons are deontically significant. In a nutshell, I will argue that explanations of why epistemic reasons are deontically significant generalize to apparent epistemic reasons. My argument relies on an epistemic version of the theory of the second best, which aims at explaining what agents should do when the best option is unavailable to them. In such cases, agents should opt for the best option *available*, or the *second-best* option. In section 2.1, I clarify why the Deontic Significance of Apparent Reasons is contentious. Specifically, I present the objections from idealized epistemology and truth-conduciveness. In section 2.2, I introduce and defend the theory of the second best, the notion of legitimate constraint and the notion of available belief-forming process. In section 2.3, I argue that, with respect to a second-best framework, we can vindicate the Deontic Significance of Apparent Reasons and respond to the aforementioned objections.

2.1. What's Wrong with Apparent Reasons?

In this chapter, I examine the deontic significance of belief-forming processes such as responding to apparent (but potentially non-factive) epistemic reasons. Before I do so, I wish to present some objections against the normativity of apparent reasons.

[51]Those who think that rationality has to do with responding to reasons in the broad sense have to explain why reasons that consist in facts about appearances are truth-conducive. Those who think that rationality has to do with responding to apparent reasons have to explain why non-factive apparent reasons are normative. As I explain in section 2.1, both challenges are similar.

2.1.1. Truth, Reasons and Apparent Reasons

The Deontic Significance of Apparent Reasons is a plausible alternative to other views, such as:

Deontic Significance of Truth. Agents are epistemically required to believe P if and only if P.

The problem with the Deontic Significance of Truth is this: it entails that evaluative norms and deontic norms are identical. That is, an agent ought to believe what he or she is better off believing. This does not make sense of the pre-theoretical assumption that, in some situations, believing P amounts to a leap of faith (even if P is true). Second, such a view makes no room for judgment suspension. That is, since P is either true or false, an agent is never epistemically permitted to withhold judgment concerning P. But this doesn't seem right. There seems to be cases in which an agent is epistemically required to withhold judgment concerning P.[52]

In view of the foregoing, responding correctly to apparent reasons one has is more plausible than the truth norm. The Deontic Significance of Apparent Reasons can explain why the deontic and the evaluative do not always coincide, and it can explain why agents are sometimes required to withhold judgment concerning P.

Still, it could be argued that there is an even better alternative, namely, responding correctly to reasons (in the narrow sense) one has. As with apparent reasons, responding to reasons one has is less demanding than a truth norm—it can explain why the deontic and the evaluative sometimes come apart, and it can make sense of the idea that, in some cases, agents ought to withhold judgment concerning P. Furthermore, by way of contrast with reasons-responsiveness, the view according to which agents ought to respond correctly to their apparent reasons faces three important challenges. The next subsections present them.

[52]As I will explain later in this chapter, we can still make sense of the Deontic Significance of Truth as an epistemic first-best (or ideal).

2.1.2. The Objection From Crazy Apparent Reasons

Here is a quick objection against the deontic force of apparent reasons. Many versions of this objection can be found in print.[53] For present purposes, I'll focus on Wedgwood's version, who says:

Crazy False Belief Objection. "Suppose that the agent insanely believes that the pattern of tea leaves at the bottom of her teacup is a divine sign confirming the truth of the proposition that she is immortal and will never die. Then according to these accounts [of responsiveness to apparent reasons], this belief—even though it is by hypothesis an insanely irrational belief—makes it the case that the agent ought to... believe that she is immortal.... Few philosophers would be willing to say... that the agent 'ought' to believe that she is immortal." (Wedgwood 2017, 58-9)

The Crazy False Belief Objection roughly says this: suppose agents should respond to their apparent reasons. Then, we are led to the absurd conclusion that insane agents should believe what is entailed by their crazy beliefs. Obviously, there is something wrong with this conclusion. So, it must be false that agents should respond to their apparent reasons.

The problem with this objection is that, if it were conclusive, it sets the bar very low for responding correctly to reasons one has. Recall that the only difference between reasons and apparent reasons is that the former are facts, while the latter are not necessarily facts. So, if Wedgwood has described a case in which the agent has correctly responded to his or her apparent reasons, we can think of a very similar case in which crazy agents respond correctly to their reasons. We simply need to assume that the agent's beliefs are true instead of false. For instance, consider the following case:

Crazy True Belief Objection. Jane insanely believes that the pattern of tea leaves at the bottom of her teacup is a divine sign confirming that she is immortal and will never die. Suppose her beliefs are true—as a matter of fact, *it is a divine sign*, and *she is immortal*. So while Jane insanely believes such propositions, she has reasons to believe that she will never die. But, as in the Crazy False Belief Objection, it is still patently clear that Jane should not believe what is entailed by her crazy beliefs.

Note that the only difference between the above objections is that, in the second one, Jane has true beliefs. So if Jane has apparent reasons in the Crazy False Belief Objection, she has reasons in the Crazy True Belief Objection. However, most philosophers will refuse to admit that, in the Crazy True Belief Objection, Jane is responding correctly to her reasons.

[53] See, e.g., Kiesewetter (2017) and Broome (2013).

Such an account of reasons-responsiveness sets the bar too low.[54] The fact that Jane has crazy (or irrational/unjustified) beliefs is an obstacle to thinking that she is reasons-responsive, even assuming that all her beliefs are true.

If that is right, then it is also false that the agent is responding to her apparent reasons in the first case. If having crazy true beliefs does not qualify as having reasons, then having crazy false beliefs does not qualify as having apparent reasons. The only difference between responding to reasons and responding to apparent reasons should be the factivity condition. Thus, the Crazy False Beliefs Objection is inconclusive.

2.1.3. The Objection From Ideal Epistemic Worlds

The first serious objection to the deontic significance of apparent reasons comes from appeals to approximation in epistemology. Philosophers have a tendency to think in terms of *ideal conditions* and *approximation* of such ideal conditions.[55] For example, it could be suggested that we ought to approximate ideal epistemic figures to a certain degree. Naturally, ideal epistemic figures are agents with knowledge or true beliefs, no false beliefs, impeccable reasoning, and so forth.[56]

Now, assume that some such view is correct. It is then hard to see the role played by apparent reasons (and, especially, of non-factive apparent reasons) in ideal epistemic scenarios. For instance, when we think of ideal epistemic agents, we do not think of massively deceived agents who believe in accordance with their hallucinations. It is patently clear that in an ideal world, agents do not respond to their non-factive apparent reasons. We can even assume that apparent non-factive reasons are *entirely ignored* by ideal agents in ideal conditions. So, if the approximation approach to epistemic norms is correct, it is hard to see the deontic significance of apparent reasons.

---

[54]Feldman (1988) and Schroeder (2011) make a similar point.

[55]Staffel (2017) calls such a view the Imitation Thesis. Smithies (2015) endorses a similar view. See also Christensen (2004, chap. 6), Pasnau (2013) or Talbott (2016, sec. 6.1.A) for discussion in various contexts. Related views can be found in argumentation theory, where idealized argumentative scenarios are used as normative profiles against which real-life argumentation is evaluated (Walton 1988, 243).

[56]See Pasnau (2013, sec. 4) for discussion.

In sum, a first objection against the deontic significance of apparent reasons could roughly be formulated as follows:

Approximation Argument. Epistemic norms bear deontic significance because they are a sufficiently good approximation of epistemic ideals. However, responding to apparent reasons (and, especially, to non-factive apparent reasons) is not part of ideal epistemic scenarios. So, responding to apparent reasons does not bear deontic significance.

2.1.4. The Objection From Epistemic Teleology

A second serious objection to the deontic significance of apparent reasons comes from epistemic teleology. According to the teleological accounts of epistemic norms, we ought to satisfy epistemic norms because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). Since truth is the final epistemic goal, we then have a goal-oriented explanation of why we ought to satisfy epistemic norms.[57] In such a perspective, epistemic norms are *conducive* to the evaluative norm of belief—that is, they are truth-conducive.

The deontic significance of apparent reasons is problematic in a teleological perspective. Consider the case in which Debby is a brain in a vat who responds to her apparent reasons. In such a case, Debby is *actually* unreliable but *counterfactually* reliable. That is, *if she found herself in the right conditions*, her ratio of true to false beliefs would be good. However, since she only responds to apparent non-factive reasons, her ratio of true to false beliefs is bad. Hence, following a teleological account of epistemic norms, it seems that Debby's belief-forming process has no positive epistemic status. Of course, in normal conditions (for example, if Debby were taken out of the vat and followed the same belief-forming processes), she would be reliable. So, it seems that a teleological explanation of the deontic significance of rationality yields the following verdict: responding to apparent reasons is merely *counterfactually* normative, and responding to reasons is genuinely normative.

So, there is a second objection against the deontic significance of apparent reasons, as in the following:

[57]There could also be mixed vindications of the deontic significance of epistemic norms (e.g., vindications that rely on teleological and nonteleological components). For the sake of simplicity, I leave this possibility aside here.

Epistemic Teleology Argument. Epistemic norms bear deontic significance because they are truth-conducive (or optimize an agent's ratio of true to false beliefs). However, responding to apparent non-factive reasons is not necessarily truth-conducive. So, responding to apparent non-factive reasons does not necessarily bear deontic significance.

In summary, vindications of the deontic significance of apparent reasons face at least two important difficulties: one comes from the approximation of ideals in epistemology, and the other comes from value-conduciveness. My argument in favour of the deontic significance of apparent reasons will provide a response to both objections.

The Epistemic Teleology Argument has been ignored in recent vindications of the normativity of epistemic rationality. Indeed, philosophers like Kiesewetter (2017) and Lord (2018) have suggested that rationality is normative because it has to do with responding to reasons one has. However, they employ the broad notion of reasons, which includes facts about appearances. As Kiesewetter indicates:

> If A's total phenomenal state supports P, and P would—if true—be an available reason for (or against) believing Q, then A's appearances provide an equally strong available reason for (or against) believing Q.... whenever some facts provide evidence for us, our appearances provide sufficient 'backup evidence' that would support the same beliefs [if we were deceived]. (Kiesewetter 2017, 173–74)

The problem is that facts about appearances are not necessarily truth-conducive. The fact that Debby is hallucinating that P might be a reason for her to believe that P, but such a putative reason does not make it *more probable* that P is true. Such a putative reason is not *truth-conducive*.[58] On the assumption that facts about mental states can count as reasons to believe P, agents who respond correctly to their epistemic reasons might not get closer to the truth at all. And recall that one of Kolodny's (2007b, 231) objection against the normativity of rationality comes from truth-conduciveness. Hence, if we are to provide a convincing vindication of the deontic significance of epistemic rationality, we need to take this objection seriously.

[58]Kiesewetter assumes that epistemic reasons for believing P make it more likely that P (Kiesewetter 2017, 168–69, 175). However, this aspect of epistemic reasons conflicts with his broad account of epistemic reasons.

2.1.5. Summary of the Problems

In summary, vindications of the normativity of apparent reasons face at least three important difficulties. First, instead of arguing that apparent reasons are deontically significant, why not simply accept that reasons are deontically significant? Second, there are two serious objections against the normativity of apparent reasons: one comes from the approximation of ideals in epistemology, and the other comes from truth-conduciveness.

My argument in favour of the normativity of apparent reasons solves these difficulties. First, I will derive the normativity of apparent reasons from the normativity of reasons. So, we can't simply argue that reasons are normative while apparent reasons aren't: the normativity of reasons and the normativity of apparent reasons rise and fall together. Second, my argument provides responses to the objection from ideal epistemic worlds and the objection from epistemic teleology.

2.2. An Epistemic Theory of the Second Best

In this section, I develop and defend the framework for an epistemic theory of the second best. This theory will provide a direct response to the Approximation Argument.

2.2.1. Best and Second-Best Belief-Forming Processes

The General Theory of the Second Best is an economic theory designed for explaining what agents can and should do if the Pareto-optimal option is unavailable. In other words, the theory of the second best allows us to determine what is the *absolute* optimal option (the first-best option) as well as the optimal option *relative to a set of constraints* (the second-best option). As Lipsey and Lancaster explain:

> Perhaps the best way to approach the problem of defining the scope of the theory of second best is to consider the role of constraints in economic theory.... in the theory of the Paretian optimum, certain constraints are assumed to be operative and the conditions necessary for the maximization of some function subject to these constraints are examined. In the theory of second best there is admitted at least one constraint additional to the ones existing in Paretian optimum theory and it is in the nature of this constraint that it prevents the satisfaction of at least one of the Paretian optimum conditions. (Lipsey and Lancaster 1956, 12)

Second-best options matter to agents. To see why, consider the following example. In a world in which people never get sick (and indeed never could get sick), there is no need for taking measures to prevent getting sick. However, in a world in which people do get sick, taking preventative measures is important. In short, when the best option of total immunity to sickness is not available, the second-best option is that of taking appropriate measures to minimize sickness. While this option is less than ideal, it clearly bears significance: taking the appropriate means to minimize sickness is significant because this is the best option available in this world. So, second-best options matter to agents.

We can adapt Lipsey and Lancaster's framework to epistemic norms and belief-forming processes.[59] In the epistemic realm, what is the *best* belief-forming process? In a teleological perspective, the best belief-forming process would be something similar to: "A believes P if and only if P is true." That is, forming the belief that P in response to P's truth is the maximally truth-conducive process. When agents follow such a belief-forming process, they end up with an infinite number of true beliefs and no false ones. Yet epistemologists often refer to other ideals that do not match this absolute epistemic ideal. The Bayesian ideal of rationality is a good example. Bayesian epistemologists often assume that idealized rational agents have unlimited cognitive capacities and reason perfectly (by way of contrast with theories of bounded rationality, where bad reasoning and limited cognitive capacities are taken into account). But they do not think that rational agents are infallible and omniscient. So, the Bayesian ideal of rationality is not the absolute best epistemic ideal. The lesson here is this: in addition to the absolute epistemic first-best, there can be some "local" first-best belief-forming processes (such as the ideal of epistemic rationality). For the moment, I focus on the absolute best belief-forming process and leave "local" first-best belief-forming processes aside.

Now, there are legitimate constraints explaining why this process is unavailable to agents like us. For instance, perhaps agents with unlimited cognitive capacities who need no guidance are epistemically required to believe P if and only if P, but lack of omniscience or the need for guidance could be *legitimate constraints* on the available belief-forming processes. In other words, there can be legitimate constraints which *exclude* processes such as

[59]I am not the first one to come up with this idea. See note 62.

"believe P if and only if P." In view of the foregoing, here is how we can define the notion of constraint on available processes:

Constraint on Available Processes. A set of legitimate constraints on processes C_A is a collection of statements which explains why some belief-forming processes are unavailable in a given context.

A quick clarificatory remark: it is possible that more than one belief-forming processes are second-best. There are at least two explanations of why this is possible. First, in some specific contexts, two belief-forming processes can warrant exactly the same attitudes, and so be equally optimal with respect to the truth norm of belief. Second, there can be equally optimal but incompatible belief-forming processes. For example, Titelbaum and Kopec (forthcoming; m.s.) have argued that there are equally reliable but incompatible standards of reasoning. So, if standard 1 is a second-best, standard 2 is incompatible with standard 1 and both standards are equally reliable (as Titelbaum and Kopec suggest), then standard 2 is also a second-best.

The above framework provides an initial adaptation of the General Theory of the Second Best to epistemic norms. The optimal unconstrained belief-forming process is an epistemic first-best. The optimal belief-forming process which satisfies the legitimate constraints on available options agents have is an epistemic second-best.

2.2.2. Response to the Approximation Argument

Second-best epistemology provides a response to the Approximation Argument. Indeed, the General Theory of the Second Best is well known for suggesting that claims concerning ideals and approximation of such ideals are ambiguous, if not misleading. Specifically, second-best options might have little or nothing in common with first-best options.[60]

Consider the following example, where second-best problems arise from the interdependence between different variables. Suppose that, in the ideal market, there are no tariffs. This is so, because free markets optimize production. Now, suppose that, in this world,

[60]Of course, approximation of ideal norms by non-ideal thinkers might be epistemically beneficial in some circumstances. See De Bona and Staffel (2018), for instance. My point here is that, in accordance with the General Theory of the Second Best, approximating an ideal is not *always* optimal in non-ideal circumstances.

every country applies tariffs. However, one country decides to adopt a free trade policy. There is a sense in which this action brings the world closer to the ideal of a free market. But does this action improve the global production? Not necessarily. In view of the global allocation of the resources, this action can be suboptimal. (Ozga 1955, 499) As Lipsey and Lancaster indicate:

> [assume] that all commodities are, in consumption, rigidly complementary, so that their production either increases or decreases simultaneously... [then] in a three country world with tariffs all around, one country may adopt a policy of free trade and, as a result, the world production of all commodities may decrease. (Lipsey and Lancaster 1956, 14)

So, the first-best scenario involves free trade for all countries, but the second-best scenario can involve tariffs in every country. A country adopting a free trade policy would approximate the ideal world, but such a policy can nevertheless be suboptimal.

In a similar spirit, here is another example which has to do with compensatory mechanisms. Suppose that the ideal world has the following two attributes: (i) people never get sick and (ii) there are no hospitals. This makes sense, because hospitals are useless if people never get sick. Now, suppose that such an ideal world is unavailable, but that such worlds are available:

Option 1. (i-) some people get sick and (ii) there are no hospitals.

Option 2. (i-) some people get sick and (ii-) there are hospitals.

Surely, if people get sick in this world, we want hospitals in this world. So, Option 2 is more valuable than Option 1. However, while Option 1 and the ideal world have one attribute in common (there are no hospitals), Option 2 and the ideal world have no attribute in common. As we can see, it would be absurd to claim that Option 1 is the best available option just because it has more attributes in common with the ideal world. Hence, the optimal (or valuable) options in non-ideal worlds might have little in common with the optimal options in an ideal world.[61]

[61]This example departs slightly from the ones discussed by Lipsey and Lancaster (1956, 26–27). Their theory aims primarily at generalizing the first kind of case I discussed (the one which has to do with tariffs). See also Daoust (m.s. e). However, the general conclusion in this example remains the same, namely, that claims in terms approximation of ideal scenarios can be misleading.

The above cases nicely illustrate the theory's main result: when at least one condition from the ideal world cannot be satisfied, approximating the other conditions might be suboptimal in the non-ideal world. This can be explained by the relationship between variables in non-ideal worlds. Or this can be explained by the fact that, in non-ideal worlds, we need to *compensate* for our imperfections: in worlds where agents get sick, it is optimal for them to have recourse to *compensatory mechanisms* such as hospitals.

The same goes for epistemic norms. Let's assume for the sake of the argument that the ideal epistemic world has the following two attributes: (i) people have unlimited cognitive capacities and (ii) they do not discriminate between pointless and significant truths (e.g., they care about all truths). However, suppose that the ideal epistemic world is unavailable, but that such worlds are available:

Epistemic Option 1. (i-) people have limited cognitive capacities and (ii) they do not discriminate between pointless and significant truths.

Epistemic Option 2. (i-) people have limited cognitive capacities and (ii-) they discriminate between pointless and significant truths (e.g., they care more about significant truths).

Surely, if people have limited cognitive capacities, we want them to use such capacities to learn significant rather than pointless truths. So, Epistemic Option 2 is superior. Yet, Epistemic Option 1 is the best approximation of the ideal epistemic world (since agents do not discriminate between pointless and significant truths in the ideal epistemic world). So, the approximation view must also be false in the epistemic realm.

With respect to the project of this chapter, such counterexamples to the approximation thesis are instructive. As I indicated above, it is patently clear that in an ideal world, agents do not respond to their non-factive apparent reasons. We can go a step further and assume that apparent non-factive reasons are *entirely ignored* by ideal agents in ideal conditions. However, this doesn't entail that, *in this world*, we should approximate the ideal world by ignoring our apparent epistemic reasons. Under the assumption that we do not have access to first-best belief-forming process, we can have recourse to radically different second-best belief-forming processes, such as responding correctly to apparent reasons one has. Hence, perhaps there is

an explanation of why, relative to the available options in non-ideal worlds, responding to apparent reasons is part of an epistemic second-best.[62]

In summary, the Approximation Argument assumes that approximation of epistemic ideals is epistemically beneficial. However, this assumption contradicts the main result of the General Theory of the Second Best. If the first-best option is unavailable to agents, it can be beneficial not to approximate ideal scenarios. This leads me to reject the Approximation Argument. Of course, one thing remains to be settled: how can we show that responding correctly to apparent reasons one has is part of an epistemic second-best? This is what I will try to show in the next section.

2.3. Responding to Apparent Reasons Is Deontically Significant

In this section, I give support to the claim that apparent reasons are deontically significant and I respond to the Epistemic Teleology Argument. My argument relies in part on the epistemic theory of the second best I sketched in the previous section, and goes as follows:

(P1) Some legitimate constraints explain why responsiveness to reasons bear deontic significance.

(P2) There are misleading reasons.

(P3) Following (P1) and (P2), responsiveness to misleading reasons bear deontic significance.

(P4) Misleading reasons and non-factive apparent reasons share the same relevant features.

(P5) Suppose that A is X (where X is a normative notion, such as "good," "bad," "required," "permitted," and so forth). Then, if A and B share the same relevant features, B is also X.

(C) Following (P1), (P3) and (P5), apparent reasons bear deontic significance.

There are three main steps in the above argument. The first step (P1) is to argue that responsiveness to epistemic reasons is part of an epistemic second-best. The second step (P2 and P3) is to argue that misleading reasons bear deontic significance. The third step (P4 and

[62]Some epistemologists are finding an interest in the theory of the second best for similar reasons. For example, according to DiPaolo (2018), norms of fallibility (which roughly states that agents ought to take their own fallibility into account when forming or revising beliefs) aims at compensating for our imperfections. DiPaolo argues that, while such norms are not part of the first-best optimal situation, they are part of the second-best optimal situation. For similar reasons, Staffel (2017) argues that non-ideal theories of epistemic rationality should not merely approximate ideal theories of rationality such as Bayesianism.

P5) is to argue that we should not make a normative distinction between misleading reasons and non-factive apparent reasons. If these steps are correct, we get the conclusion that apparent reasons (and, especially, non-factive apparent reasons) are deontically significant.

2.3.1. Defending P1

It follows from the Datum that responding to reasons is deontically significant. Given the epistemic theory of the second best, an epistemic norm or belief-forming process is deontically significant if it is part of an epistemic first-best or an epistemic second-best. So, either responding to epistemic reasons is part of an epistemic first-best, or it is part of an epistemic second-best.

Now, it should be noted that responding to reasons is not the absolute best belief-forming process agents can satisfy. As Blanshard rightly stresses, we take responding to reasons as an optimal process because we lack the ability to respond to the truth directly. According to him:

> 'Surely the only possible rule', one may say, 'is to believe what is true and disbelieve what is false.' And of course that would be the rule if we were in a position to know what was true and what false. But the whole difficulty arises from the fact that we do not and often cannot. What is to guide us then?… believe no more, but also no less, than what the evidence warrants. (Blanshard 1974, 410–11)

Here, the upshot is that the optimality of responding to reasons already presupposes some constraints on available options. Responding correctly to reasons one has can lead one to withhold judgment on whether P, or even to falsely believe P. This is less than ideal. In the best epistemic scenario, agents have all the true beliefs they can have and no false belief. So, they do not withhold judgment on whether P, or falsely believe that P. Clearly, there are better candidates for a first-best belief-forming process, such as "A believes P if and only if P is true." Accordingly, responsiveness to reasons is a second-best option. Reasons-responsiveness have deontic force with respect to the legitimate constraints agents face.

The nature of such constraints is contentious. Plausibly, logical impossibilities are a source of legitimate constraints. But what about the laws of nature or the social laws?[63] What about extremely improbable events (provided that we have a non-arbitrary definition of the notion of extreme improbability)?[64] What about a conjunction of probable events which, taken together, are extremely improbable?[65] What if it is possible to overcome a constraint, but only through a demanding course of action? I do not wish to answer these questions here. The notion of legitimate constraint is difficult to theorize. Without making an exhaustive list of possible or plausible constraints, here are some factors that could plausibly count as legitimate constraints on the set of available belief-forming processes.

First, there could be proposition-related constraints. The nature of some propositions can explain why some belief-forming processes are unavailable. For example, there are "blindspot propositions", namely, propositions that can't be truly believed. Such propositions can explain why ideal belief-forming processes (such as believe P if and only if P) are sometimes unavailable. Relative to such propositions, either one forms a false belief concerning P or withholds judgment concerning P, but truly believing P is impossible.[66]

Second and more importantly, there could be agent-related constraints. For example, some theories of epistemic rationality are concerned with credences, whereas others are concerned with full beliefs. An explanation of why some theorists privilege belief over credence is this: entertaining fine-grained degrees of belief is *too demanding* of epistemically rational agents.[67] So, there can be constraints on the type of doxastic attitudes agents can entertain. In a similar vein, it has been suggested that cognitive limitations or abilities,[68]

[63]Stemplowska and Swift (2012), Estlund (2014) and Gaus (2016, 30–33) analyze this problem for ideal theories of justice.
[64]See Gaus (2016, 30–31).
[65]As Gaus indicates, "suppose that a theory posits five soft constraints each of which yields a 20 percent chance that the utopia will arise, and all come to bear. Given the joint probabilities, there is only about 0.00032 chance utopia will come about" (2016, 34).
[66]See Bykvist and Hattiangadi (2007), Olinder (2012) and Raleigh (2013) on blindspot propositions. Relatedly, there can be safespot propositions, namely, propositions that are "guaranteed to be true provided the subject adopts a certain attitude towards it" (Raleigh 2015, 309). See also Kopec (2015) and Reisner (2007; 2014) on safespots.
[67]See Christensen (2004, 144–50) for discussion.
[68]Among others, see Lord (2015), Paul and Quiggin (2018), Todd and Gigerenzer (2000), Loftus (2005), McCain (2008), Podgorski (2016) or Smithies (2015; 2016).

fallibility,[69] the need for guidance,[70] doxastic involuntarism[71] or the accessibility of reasons[72] are legitimate constraints on the available belief-forming processes.

Third, there can be constraints of 'short-term optimality', as in the following case:

Deal with God. God offers you the following deal: if you form beliefs randomly for fifty years, you will afterwards be given the opportunity of being omniscient and your human cognitive limitations such as the need for guidance will disappear. That is, the absolute best belief-forming process (say, "believe P if and only if P") will be available to you, but only after forming beliefs randomly for fifty years.

The Deal with God is a special case in which, in order to access the perfect belief-forming process, you must first satisfy a suboptimal belief-forming process. Specifically, this is a case in which short-term optimality conflicts with long-term optimality. In such a case, the need for short-term optimality could constitute a legitimate constraint excluding long-term optimality.[73]

As we can see, there could be numerous legitimate constraints on available belief-forming processes, and offering a satisfactory account of such constraints goes beyond the project of this chapter. Most (if not all) of them are contentious. The good news is that, in the remainder of this chapter, we won't need to pinpoint the legitimate constraints on the set of available belief-forming processes. *By hypothesis*, if responding to reasons is part of an epistemic second-best, there is a set of legitimate constraints C_A which explains such a fact. This is all we need to know for the moment.

2.3.2. Defending P2 and P3

P2 states that there are misleading reasons. Many simple examples support this claim. For instance, it can be a fact (or a true proposition) that the combination 0-49-21-34-38-11 had 0.00000001 chance to be tonight's winning lottery combination, and such a fact supports the

[69]See DiPaolo (2018).
[70]See, for example, McHugh and Way (2017, 131–32) or Reisner (2009, 244–46, 248–50). See Way and Whiting (2016) for discussion of the relationship between the need for guidance and the "Ought-Implies-Can" principle.
[71]See notably Owens (2002) for discussion.
[72]See Wedgwood (2002) and Hatcher (2016) for discussion.
[73]Except perhaps Foley's (1987, 8–14) distinction between our current and long-term epistemic goals, I am not aware of anyone discussing such cases in the context of epistemic teleology. In political contexts, this is known as the problem of "rugged landscapes." See Gaus (2016, sec. II.2). The separateness of propositions could also explain why agents ought not to take the Deal with God. See Berker (2013a).

conclusion that 0-49-21-34-38-11 is a losing combination. However, such a fact can be misleading: for instance, it is possible that 0-49-21-34-38-11 is the winning combination. In other words, this fact points in the wrong direction. It will draw a reasons-responsive agent away from the truth. Therefore, there are misleading reasons.

Of course, one could be an infallibilist about reasons and argue that beliefs supported by reasons are necessarily true. So, the fact that the combination 0-49-21-34-38-11 had 0.00000001 chance to be tonight's winning lottery combination might not *support* the belief that the ticket is a loser. Granted, but the same problem would arise with respect to other doxastic states such as credences. Plausibly, if I know that the combination 0-49-21-34-38-11 had 0.00000001 chance to be tonight's winning lottery combination, then I should entertain a credence of 0.00000001 in "0-49-21-34-38-11 is the winning combination."[74] However, such a credence is very far from the epistemically ideal credence of 1 in "0-49-21-34-38-11 is the winning combination." So, the fact that 0-49-21-34-38-11 had 0.00000001 chance to be tonight's winning lottery combination is a misleading reason, in the sense that it leads one to entertain highly inaccurate credences. So again, the conclusion that there are misleading epistemic reasons is correct.

As for P3, it is a direct consequence of P1 and P2. If reasons bear deontic significance and there are misleading reasons, misleading reasons bear deontic significance.

2.3.3. Defending P4 and P5

P4 is a central step in the argument. Misleading reasons and non-factive apparent reasons share the same normatively relevant features. They are outweighed (or defeated) as soon as agents recognize them, they are undermined relative to the goals of believing what's true and not believing what's false, they make knowledge unsafe and we are typically unable to distinguish them from non-misleading reasons. In such a context, there is no ground for making a normatively relevant distinction between them.

[74]At least assuming that, if agents respond correctly to their reasons, their credences reflect their knowledge of the objective probabilities.

A good indication of the fact that misleading reasons and apparent reasons share the same normatively relevant features comes from the *reaction* of rational agents to discovering that they have misleading reasons or non-factive apparent reasons. It is rational for agents to respond to merely apparent reasons only insofar as they are *not aware* that the considerations they are responding to are merely apparent. If Debby, the brain in a vat, knew a procedure for identifying non-factive reasons, she would discard such considerations immediately, which means that they would merely count as *defeated* apparent reasons to her. Consider the following case:

Refraction of Light. Debby sees a stick that is thrust partly into the water. The stick appears to be bent to her. However, Debby is aware that this is a mere optical illusion caused by the refraction of light. For that reason, Debby ignores her misleading impression that the stick is bent.

As we can see in the above case, Debby has a defeated apparent reason to believe that the stick is bent. It appears to her that she perceives that the stick is bent, but she is aware that such a perception is merely apparent and does not capture any fact. This is why she discards her apparent perception that the stick is bent—such an apparent perception doesn't really count as an apparent reason to her. Relatedly, if Debby knew that she was a brain in a vat and that all of her perceptions are merely apparent, she would not rationally take the content of her perceptions as reasons to believe anything. These examples and others like it suggest that apparent reasons remain undefeated only if they are indistinguishable from (genuine) reasons.

Responding to apparent but *non-factive* reasons is no different than responding to *misleading* reasons. If rational agents know that X is a misleading reason to believe P, they take X to be a *defeated* reason. The only explanation of why rational agents respond to misleading reasons is their unawareness that such considerations are misleading.[75] For example, imagine that Debby is informed that the combination 0-49-21-34-38-11 had 0.00000001 chance to be tonight's winning lottery combination, and that she has no other information concerning tonight's winning lottery combination. Debby will form beliefs in accordance with such an information only if she is unaware that the ticket is, in fact, a winner.

[75]Horowitz (2014b, 43) makes a similar observation.

This is why an epistemically rational agent would prefer to discard misleading reasons. Indeed, imagine that Debby knows that 0-49-21-34-38-11 is tonight's winning lottery combination. In such a case, she would discard the information that the combination had 0.00000001 chance to be tonight's winning lottery combination. So, as with apparent non-factive reasons, misleading reasons are undefeated only insofar as an agent cannot distinguish them from non-misleading reasons.

In summary, when rational agents discover that they have misleading reasons or non-factive apparent reasons, they react in the same way: they discard these considerations, or they take such considerations to be defeated. This confirms P4, which states that non-factive apparent reasons and misleading reasons share the same normatively relevant features.

P5 is the consequence of a plausible supervenience principle analogous to the ones found in ethics. Ethical supervenience principles state that the normative properties of X supervene on the relevant non-normative properties of X. For instance, suppose that two stacks of hay share the same normatively relevant features—they are equally big, tasty, reachable, and so forth. Of course, the stacks of hay might not be identical: one might be 10 meters on A's left, while the other might be ten meters on A's right. Accordingly, there can be some differences between these stacks of hay. But we can assume that there is no normatively relevant difference between them, in the sense that the considerations favouring one will also favour the other, and vice versa. The fact that one stack is on A's left rather than on A's right is normatively irrelevant. We can't make a normatively relevant distinction between these stacks of hay. If they share the same normatively relevant features, no normative distinction should be made between them.

2.3.4. Back to the Deontic Significance of Apparent Reasons

If the Datum and P1 to P5 are correct, the deontic significance of apparent reasons is vindicated. That is, we can't conclude that reasons are deontically significant without also concluding that apparent reasons are deontically significant.

Here is a simpler way to understand the argument I put forth in the above sections. To begin with, consider the following belief-forming processes:

(Pr-α) A believes P if P is sufficiently supported by A's non-misleading reasons;

(Pr-β) A believes P if P is sufficiently supported by A's reasons;

(Pr-γ) A believes P if P is sufficiently supported by A's apparent (factive or non-factive) reasons.

Suppose that the Datum is correct, but that only reasons bear deontic significance. Clearly, responding to reasons is not an epistemic first-best (an epistemic first-best would probably look like "A believes P if and only if P"). Under such assumptions, Pr-β is would be a deontically significant second-best. But if Pr-β is a second-best, there has to be an explanation of why Pr-α is unavailable. Indeed, responding to non-misleading reasons is equally or more truth-conducive than responding to reasons, since the former leaves out misleading considerations while the latter does not. So, at every possible world, Pr-α is equal or superior to Pr-β. In view of the foregoing, there has to be a legitimate constraint excluding Pr-α from the set of available belief-forming processes. Otherwise, Pr-β would not be an epistemic second-best (there would be an equal or superior option available at every possible world).

Now, think of your favourite legitimate constraint (call it C_A) explaining why Pr-α is unavailable. For instance, suppose that you think agents lack the ability to distinguish misleading reasons from non-misleading reasons, and this explains why Pr-α is unavailable. Here is the problem: misleading reasons and non-factive apparent reasons share the same normatively relevant features. So, since C_A explains why agents are permitted to respond to their misleading reasons, it also explain why they are permitted to respond to their apparent non-factive reasons. Otherwise, C_A makes an illegitimate distinction between misleading reasons and apparent non-factive reasons. Hence, it can't be right that only reasons bear deontic significance: if C_A explains why misleading reasons bear deontic significance, it will end up explaining why apparent reasons also bear deontic significance.[76]

[76]One could go a step further and argue that Pr-β cannot be an epistemic second-best, because it is either unavailable or suboptimal relative to the processes available to one. A purported argument for this conclusion runs as follows: in cases where Pr-α is available, Pr-β is suboptimal, and in cases where some constraints exclude Pr-α, the constraints also exclude Pr-β. The problem with such a line of reasoning is that we can imagine possible worlds where Pr-α and Pr-β are equally optimal, and so are both second-best. For instance, in possible worlds where agents never possess misleading reasons (or non-factive apparent reasons), Pr-α, Pr-β and Pr-γ are equally optimal. So, Pr-β could be an epistemic second-best. Yet in accordance with these remarks, perhaps one could argue that Pr-β can't be the *unique* second-best process.

2.3.5. What About the Epistemic Teleology Argument?

It could be argued that the above argument does not provide a satisfactory answer to the Epistemic Teleology Argument. Indeed, responding to reasons is more truth-conducive than responding to apparent (factive or non-factive) reasons. For instance, a brain in a vat like Debby who responds to non-factive apparent reasons has a very bad ratio of true to false beliefs. Hence, since the process of responding to apparent reasons is less truth-conducive than the process of responding to reasons, it seems that responding to apparent reasons does not bear deontic significance in a teleological perspective. So, how does the above argument answer the Epistemic Teleology Argument?

A second-best framework like the one I have developed will simply tell us to take the legitimate constraints on available belief-forming processes into account. If Debby can't identify apparent non-factive reasons, she satisfies the best (or most truth-conducive) *available* belief-forming process in responding to her apparent reasons. Naturally, since Debby is a brain in a vat, responding to apparent reasons does not secure a good ratio of true to false beliefs for her. Nevertheless, Debby is still doing the best she can. It's a fact of our epistemic lives that, in some specific situations, doing our best won't secure a good ratio of true to false beliefs.

The above reasoning is also correct for agents who respond to their reasons. Even if agents respond to their reasons, they might end up with a bad ratio of true to false beliefs. Indeed, recall that there are misleading reasons. So, there could be cases where an agent is very unlucky and only has misleading reasons. In such a case, an epistemically rational agent's ratio of true to false beliefs won't be good. Again, it's a fact of our epistemic lives that, even if we do our best to acquire true beliefs and avoid false beliefs, we might still end up with a bad ratio of true to false beliefs.

Of course, it could be argued that, in *normal circumstances*, responding correctly to reasons is truth-conducive. However, in normal circumstances, responding correctly to apparent reasons is also truth-conducive. So, we can't make a distinction between responding to apparent reasons and responding to reasons by appealing to what's going on in normal circumstances.

2.4. Conclusion

2.4.1. Main Results of This Chapter

In this chapter, I argued that the Deontic Significance of Apparent Reasons is correct. This also confirms the Minimal Normative Hypothesis, according to which responding correctly to (apparent) reasons one has is deontically significant (see section 1.4.1). Indeed, the deontic significance of reasons is almost self-explanatory—that is, reasons to believe P are facts counting in favour of believing P, and I assume that the "counting in favour of" relation is normative. Following the argument of this chapter, apparent reasons (or reasons in the broad sense) are also deontically significant. Thus, responding correctly to reasons and apparent reasons one has is deontically significant.

In addition to solving the dispute surrounding the Deontic Significance of Apparent Reasons, this chapter offers an alternative to existing theories of the relationship between ideal and non-ideal epistemic norms. Typically, philosophers assume that norms in non-ideal worlds consist in sufficiently good approximations of norms in ideal worlds. I have argued that such a view is misleading. This explains why believing in accordance with apparent reasons one has can be permitted or required in non-ideal worlds, even if agents never respond to apparent non-factive reasons in ideal worlds.

The chapter also answers a worry put forth by Niko Kolodny. According to him, a reason why epistemic rationality is not normative is that being rational does not guarantee a good ratio of true to false beliefs (see section 1.2.5). This leads Kolodny to think that epistemic reasons in the narrow sense are normative, but not apparent reasons. My response to him is this: of course, if one is a deceived brain in a vat, responding to apparent reasons one has won't secure a good ratio of true to false beliefs. But neither does responding to *misleading* epistemic reasons in bad circumstances. It's a fact of our epistemic lives that, whether we respond to apparent reasons or to reasons in the narrow sense, we can end up with a bad ratio of true to false beliefs. Insofar as reasons in the narrow sense can be misleading, turning to epistemic reasons and giving up apparent reasons won't change this fact.

In view of the foregoing, if Kolodny's objection from truth-conduciveness were correct, epistemic reasons would not be normative. But this contradicts the Datum (and his own claim that epistemic reasons are normative). So, his objection is inconclusive by his own lights.

2.4.2 The Elimination of Structural Requirements of Rationality

Even if the Minimal Normative Hypothesis is correct, we still need to determine whether coherence requirements are an obstacle for vindicating the normativity of epistemic rationality. In the next chapter, I begin the examination of coherence requirements of rationality. Chapter 3 is devoted to Inter-Level Coherence and Intra-Level Coherence. Chapters 4 and 5 are devoted to Consistency.

Chapter 3. Intra-Level Coherence, Inter-Level Coherence and Epistemic Reasons

Chapter summary. This chapter addresses the status of coherence requirements of rationality (and, in particular, the compatibility between such requirements and reasons-responsiveness). It seems that epistemically rational agents should avoid incoherent combinations of beliefs and should respond correctly to their epistemic reasons. However, some situations seem to indicate that such requirements cannot be simultaneously satisfied. In such contexts, assuming that there is no unsolvable dilemma of epistemic rationality, either (i) it could be rational that one's higher-order attitudes do not align with one's first-order attitudes or (ii) requirements such as responding correctly to epistemic reasons that agents have are not genuine rationality requirements. This result doesn't square well with plausible theoretical assumptions concerning epistemic rationality. So, how do we solve this puzzle? In this chapter, I will suggest that an agent can always reason from infallible higher-order reasons. This provides a partial solution to the above puzzle.

Meet Doctor Watson, Sherlock Holmes's assistant. While he rarely matches Holmes's reasoning skills, Watson is an epistemically rational reasoner.[77] Now, imagine that Watson finds himself in the following situations:

Clear Evidence. Watson has sufficient evidence of numerous distinctive features X (the type of murder, the type of victim, the crime scene's location, etc.). Given features X, Watson can deduce that the killer is Jack the Ripper.

Fallible Reasons. Watson analyzes numerous distinctive features X (the type of murder, the type of victim, the crime scene's location, etc.). He finds a justificatory chain leading to the conclusion that the killer is Jack the Ripper. However, he is aware that the reasons he responded to are fallible to a certain degree.

Bad Reasoning. Watson concludes that the killer is Jack the Ripper on the basis of numerous distinctive features X (the type of murder, the type of victim, the crime scene's location, etc.). However, he also has evidence (i) that Holmes thinks that he (Watson) made a mistake in processing the evidence and (ii) that Holmes is almost always reliable. For example, Holmes could suggest that, on that particular occasion, Watson reached a conclusion through incorrect reasoning.

[77] I borrowed these "Watson cases" from Coates (2012) and Horowitz (2014a).

Let's assume that, in cases like Clear Evidence, Watson is epistemically rational in concluding that Jack the Ripper is the killer. However, as I will explain below, things get complicated in cases like Bad Reasoning or Fallible Reasons. In such cases, it isn't clear how Watson will rationally weight the evidence he has or evaluate his own reasoning.

Several authors have recently suggested that, in cases like Bad Reasoning or Fallible Reasons, it is rational for Watson to hold an akratic combination of attitudes (Coates 2012; Lasonen-Aarnio 2014; Lasonen-Aarnio m.s.). Others have suggested that such cases show that responding to epistemic reasons is not a genuine requirement of epistemic rationality, or at least that responding to epistemic reasons can conflict with coherence requirements (Worsnip 2015). Let's call this a Rational Puzzle:

Rational Puzzle. At least one of the following verdicts is correct: (i) epistemic akrasia can be rational, or (ii) requirements such as responding correctly to epistemic reasons are not genuine rationality requirements.

Rational Puzzle is problematic because it does not cope well with plausible assumptions concerning epistemic rationality. In particular, it is hard to imagine that an epistemically rational agent sometimes has to choose between responding correctly to his or her reasons and maintaining internal coherence.

In this chapter, I shed light on the above puzzle. First, it is sometimes helpful to determine that what appears to be a new problem is, in fact, very similar to a well-known one. I will suggest that Rational Puzzle is essentially related to traditional problems of responding to fallible reasons such as the lottery paradox. Specifically, if the fallibilist solution to the lottery paradox is correct, then it could be rational for an agent to hold an akratic combination of attitudes. Nevertheless, I will suggest that an agent never has to choose between responding to his or her reasons and avoiding akratic combinations of attitudes, because he or she is always in a position to satisfy both requirements.

In section 3.1, I will present Rational Puzzle and explain why cases like Bad Reasoning or Fallible Reasons are closely related to this puzzle. In section 3.2, I will argue that Rational Puzzle holds only if a rational agent can have sufficient epistemic reason to believe that "he or

she has sufficient epistemic reason to believe P," while having sufficient epistemic reason against believing P. I will then explain that such situations are possible only if higher-order epistemic reasons are sometimes fallible.

This will lead me, in section 3.3 and 3.4, to analyze the possibility of fallible higher-order epistemic reasons. I will argue that, while there can be fallible higher-order epistemic reasons, an agent can *always* respond to infallible higher-order epistemic reasons. Furthermore, relative to rational reasoning, responding to infallible higher-order epistemic reasons appears to be preferable. In other words, I will argue that a rational agent would prefer responding to infallible higher-order reasons. This provides a partial solution to Rational Puzzle: while this chapter does not rule out the possibility of rational epistemic akrasia, (i) no epistemically rational agent is required to maintain such a combination of attitudes and (ii) remaining in such a state seems undesirable. In conclusion, I discuss the implications of my arguments in the debate on the normativity of epistemic rationality.

3.1. Rational Believers, Enkratic Requirement(s) and Rational Puzzle

3.1.1. Some Quick Reminders

An ideally rational agent satisfies all rationality requirements. As discussed in section 1.1.2, structural requirements govern relations among multiple attitudes. They are, for the most part, coherence requirements. Here are two putative coherence requirements of rationality:[78]

Consistency. Epistemic rationality requires that, if A believes that P, then it is false that A believes that ~P.

Intra-Level Coherence. Epistemic rationality requires that, if A believes that P_1, believes that P_2, ..., and believes that P_n, then it is false that A believes that $\sim(P_1 \wedge P_2 \ldots \wedge P_n)$.

Consistency is logically weaker than Intra-Level Coherence. For example, simultaneously believing P, believing Q and believing ~(P^Q) violates Intra-Level Coherence but such a combination of beliefs does not necessarily violate Consistency. For the moment, I

[78]See notably Broome (2005, 322; 2007a, 355; 2013, sec. 9.2).

will only assume that Consistency is correct, and I will come back to Intra-Level Coherence in section 3.3 when discussing lottery cases.

Substantive requirements govern how agents form and revise their attitudes in accordance with their epistemic reasons (for the sake of simplicity, I'll here refer to reasons in the broad sense, which include facts about appearances). For example, when an agent has sufficient epistemic reason to believe P, this seems to put him or her under a normative pressure to come to believe P. In other words, rationality requires that, if A has sufficient epistemic reason (understood in the broad or in the narrow sense) to believe P, A believes that P.

A brief reminder: as I indicated in section 1.5.3, I here assume that epistemic reasons are represented by *epistemic probabilities*, understood as the probabilities warranted by an agent's body of epistemic reasons. Fallible reasons to believe P warrant an epistemic probability of less than 1 in P, and infallible reasons to believe P warrant an epistemic probability of 1 in P. Also, while rational credences are *not identical* to epistemic probabilities, they *track* epistemic probabilities. For example, if P's epistemic probability is 0.9 relative to a body of epistemic reasons, then it is rational for an agent who has such a body of epistemic reasons to entertain a credence of 0.9 in P.

3.1.2. Inter-Level Coherence (or Reasons Enkrasia)

Akratic agents seem to be irrational. Many people have suggested that akrasia reveals inter-level incoherence—that is, incoherence between an agent's first and higher-order attitudes.[79] The "anti-akrasia constraint" can be defined as follows:

[79]Alexander (2013) suggests that, when agents have a higher-order doubt about P, they should not take a higher-order attitude towards P. Broome (2013, 22–23, 170–71) roughly suggests that, in practical cases, failure to conform to the Enkratic requirement is an internal failure, a failure with respect to your own deliberation and standards. However, he suggests that the epistemic version of Enkrasia brings more difficulties (Broome 2013, 170-72, 216-19). D. Greco (2014) argues that epistemic akrasia leads to a kind of fragmentation or irrational inner conflict. Hinchman (2013) defends the claim that epistemically akratic agents end up in a situation of self-mistrust. According to Horowitz (2014a), epistemically akratic combinations of attitudes lead to patently bad reasoning. Reisner (2013) suggests that, while the enkratic requirement is not a rationality requirement, it is strongly connected with agentivity. According to Titelbaum, mistakes concerning rationality requirements are necessarily irrational, which implies that "no situation rationally permits any overall state containing both an attitude A and the belief that A is rationally forbidden in one's current situation" (2015, 261). Titelbaum's argument is premised on the assumption that akrasia is irrational. See also Littlejohn (2015), who endorses Titelbaum's view and adds that inter-level incoherence is the sign of an opaque mindset.

Inter-Level Coherence (or Reasons Enkrasia). Epistemic rationality requires that (if A believes that he or she has *sufficient epistemic reason* to believe P, then A believes that P).

However, we find many variants of this thesis in the literature, as in the following:[80]

Evidence Enkrasia. Rationality requires that (if A believes that his or her *evidence sufficiently supports* the belief that P, then A believes that P).

Ought Enkrasia. Rationality requires that (if A believes that he or she *ought* to believe that P, then A believes that P).

Justification Enkrasia. Rationality requires that (if A believes that he or she is *epistemically justified* in believing that P, then A believes that P).

"Rational" Enkrasia. Rationality requires that (if A believes that *rationality* requires of him or her to believe that P, then A believes that P)

Obviously, claims concerning epistemic rationality, knowledge, justification, epistemic obligations and evidence are related to epistemic reasons in *some* ways. However, we cannot assume that all the above claims are equivalent. Since I will not assume that claims concerning justification, rationality, obligations, epistemic reasons and evidence are equivalent, I will focus on Inter-Level Coherence and leave the other variants behind.[81]

Historically, philosophers have been concerned with the possibility of holding an akratic combination of attitudes.[82] More recently, philosophers have focused on the normative

Finally, many philosophers defend the claim that akrasia is similar to Moore-paradoxical doxastic states—some deeply incoherent combinations of attitudes. See notably Chislenko (2014), Feldman (2005b), Huemer (2007) and Smithies (2012).

[80]For example, Horowitz (2014a) analyzes the converse of Evidence Enkrasia, Broome (2013) considers Ought Enkrasia, Feldman (2005b) is concerned with both Justification Enkrasia, and Lasonen-Aarnio (2015) addresses "Rational" Enkrasia. Also, some putative requirements of rationality like the "RR principle" of the Fixed Point thesis are very close to "Rational" Enkrasia. See notably Conee (2010, sec. 3), Lasonen-Aarnio (m.s., sect. II), Littlejohn (2015, 5) and Titelbaum (2015).

It should also be noted that many philosophers are concerned with the oddity of combination of attitudes like the following: "P, but it is false that my epistemic reasons sufficiently support the conclusion that P" (see Horowitz 2014 on this case and see Lasonen-Aarnio m.s. for discussion). I am not convinced that this variant of epistemic akrasia is necessarily irrational. There could be cases where an epistemically rational agent believes P while believing that his or her epistemic reasons do not sufficiently support P. For example, one could be in an epistemically permissive situation where, relative to a body of evidence, incompatible doxastic attitudes towards P are rationally permitted (see notably White (2014) and Kelly (2014) on epistemic permissiveness). To avoid the debate surrounding permissiveness, the only counterexamples to Inter-Level Coherence I will consider look like the following: "I don't believe that P, but my epistemic reasons sufficiently support the conclusion that P."

[81]This generates a methodological difficulty, since the enkratic requirements discussed in the literature take distinct incompatible forms. Nevertheless, as long as it does not lead to straightforward nonsensical results, I will engage with the literature *as if* other authors had discussed Inter-Level Coherence.

[82]See notably Davidson (1982, 302-304), Mele (1988, chap. 2-3), Pears (1984, chap. 9), Ribeiro (2011) and Zheng (2001).

issue of whether an epistemically akratic combination of attitudes can be rational. These two issues are related. If agents cannot hold an akratic combination of attitudes, determining whether an epistemically akratic combination of attitudes can be rational seems pointless, since such a situation can never happen. In this chapter, I will assume that akratic combinations of attitudes are possible. I will focus on whether such combinations of attitudes are necessarily irrational.

3.1.3. The case for Rational Puzzle

I now wish to explain the case for Rational Puzzle, which is a reconstruction from two distinct positions that can be found in the literature. Since these two stands were developed independently from each other, I want to explain why these positions, taken together, constitute a puzzle.

First, suppose that there are substantive requirements of rationality, such as responding to the epistemic reasons agents have, and that there is no unsolvable dilemma of rationality. In cases like Fallible Reasons and Bad Reasoning, it could be suggested that one way to respond correctly to the evidence agents have is to transgress Inter-Level Coherence and be akratic. According to Allen Coates, if Holmes tells Watson that he is irrational in concluding that Jack the Ripper is the killer, Watson's rational response to such higher-order evidence is to believe that his epistemic reasons (including deductive reasoning and evidence) do not support the conclusion that Jack the Ripper is guilty. Similarly, suppose Watson realizes that the justificatory chain he identified for the conclusion that the killer is Jack the Ripper is supported by fallible reasons. Then, it could be rational for him to believe that his epistemic reasons do not support the conclusion that Jack the Ripper is guilty. However, recall that there are rational false beliefs. So, perhaps Watson is rational in concluding that Jack the Ripper is the killer. In such a case, Watson could be rational in believing that Jack the Ripper is guilty *and* respond correctly to his evidence in concluding that his epistemic reasons do not support that conclusion (Coates 2012, 113–15). According to Coates:

> Before he spoke to Holmes, Watson's belief was, by hypothesis, perfectly rational. And the only change in his epistemic circumstances is that he has heard Holmes's

> assessment. So any objection which claims that his belief is irrational must show that Holmes's assessment of it somehow explains why it is irrational (Coates 2012, 115).

Therefore, Watson could be rational in having an akratic combination of attitudes. Now, what about the fact that being akratic is deeply incoherent? According to Maria Lasonen-Aarnio, when an agent has higher-order evidence concerning his or her own rationality, it is not always possible to identify a single coherent combination of attitudes that he or she could hold (Lasonen-Aarnio 2014; Lasonen-Aarnio m.s.). For example, she argues that "recommending that one believe that a rule is flawed is not tantamount to recommending that one stop following the rule. That one should believe that one shouldn't φ doesn't entail that one shouldn't φ" (Lasonen-Aarnio 2014, 343). In accordance with Coates, Lasonen-Aarnio concludes that it is sometimes rational for an agent to maintain incoherent combinations of beliefs, and thus to transgress Inter-Level Coherence.

Alex Worsnip also agrees that, in some situations, Watson's evidence can support (i) that Jack the Ripper is the killer and also support (ii) that his evidence does not support that conclusion. What Worsnip rejects is that responding correctly to the evidence agents have is rationally required. Indeed, while Coates and Lasonen-Aarnio's conclusion presupposes that responding correctly to the evidence agents have is a requirement of rationality, Worsnip denies that if Watson's evidence supports P, then rationality requires of Watson that he believes that P, especially in cases where this means having an incoherent combination of attitudes. According to him, evidence-responsiveness and inter-level coherence "are, properly understood, fundamentally different kinds of normative claim, such that they should not be stated using the same normative concept" (Worsnip 2015, 6). As I indicated in the previous section, for the sake of comparability between arguments found in the literature, I'll reinterpret Worsnip's claim in terms of epistemic reasons. A plausible reinterpretation of Worsnip's conclusion is to deny that reasons-responsiveness necessarily has to do with rationality.[83]

[83]Strictly speaking, Worsnip never said such a thing. However, this strikes me as a plausible consequence of his view, since he associates coherence with rationality and argues that reasons-responsiveness is best captured by different normative claims. In view of the foregoing, it seems that reasons-responsiveness would be best captured by claims outside the realm of rationality. Also Worsnip's view is compatible with the claim that reasons-responsiveness is a source of normative pressure on agents, but such a normative pressure would not come from rationality. See also Worsnip (2016).

In summary, it seems that we must accept the puzzle. On the one hand, we can admit that reasons-responsiveness is a requirement of rationality and that there is no dilemma of epistemic rationality, but then we must give up Inter-Level Coherence. On the other hand, we can admit that there is no dilemma of epistemic rationality and that Inter-Level Coherence is a rationality requirement, but then we must give up reasons-responsiveness. Rational Puzzle seriously affects how rationality is canonically understood. *Contra* Lasonen-Aarnio and Coates, it is plausible that coherence requirements are genuine requirements of rationality, including coherence between an agent's first and higher-order attitudes.[84] *Contra* Worsnip, it seems that epistemic rationality has to do with more than mere coherence. Otherwise, if conspiracy theorists and hard-core skeptics are fully coherent, they would also be fully rational, and that doesn't seem correct.[85] *A priori*, no position is comfortable or copes well with other plausible theoretical assumptions regarding epistemic rationality.[86]

3.2. Rational Puzzle and Level-Splitting

In this section, I will argue that Rational Puzzle holds only if an agent can have sufficient epistemic reason to believe that "he or she has sufficient epistemic reason to believe P," while not having sufficient epistemic reason to believe P.[87] I will refer to these situations as cases of level-splitting.

[84]See Broome (2013, chap. 9) or Gibbons (2013, 229–34). See also note 79.
[85]See Dogramaci and Horowitz (2016) and Horowitz (2014b).
[86]A third possibility would be to maintain Inter-Level Coherence and reasons-responsiveness requirements, but to conclude that, in some situations, agents will necessarily defy the ideals of epistemic rationality. If Watson concludes that he cannot rationally respond to his epistemic reasons, he could withhold judgment on whether Jack the Ripper is guilty. However, he has sufficient evidence that Jack the Ripper is the killer, which means that he does not respond correctly to the evidence he has. But if he believes that he can rationally respond to his epistemic reasons, Watson does not respond correctly to Holmes's testimony that he is currently unable to respond to his epistemic reasons. According to David Christensen, in such a case, regardless of how Watson respond to his evidence, he could be "doomed to fall short of the rational ideal" (Christensen 2010, 212). Such a claim is controversial. Chang (2001) and Bélanger (2011) argue that all normative dilemmas can be solved. Plausibly, if rationality is supposed to offer guidance, or to consistently determine an agent's permissions and obligations, then every apparent dilemma of rationality should be solvable. This is why I here assume that putative dilemmas between Inter-Level Coherence and reasons-responsiveness are solvable. On the other hand, Hughes (2017), Sinnott-Armstrong (1996) and Williams (1965) defend the claim that there are unsolvable normative dilemmas.
[87]Lasonen-Aarnio (m.s.), Worsnip (2015) and Horowitz (2014a) reach similar conclusions.

A key feature of Rational Puzzle is that reasons-responsiveness and Inter-Level Coherence sometimes lead to incompatible verdicts. As long as higher-order epistemic reasons are coherent with first-order epistemic reasons, Inter-Level Coherence and reasons-responsiveness are compatible. For example, suppose that an agent has sufficient epistemic reason to believe that "he or she has sufficient epistemic reason to believe P" and sufficient epistemic reason to believe P. In such a case, reasons-responsiveness requires that agent to believe that he or she has sufficient epistemic reason to believe P and to believe P. Such a combination of attitudes satisfies Inter-Level Coherence. So, if an agent's first and higher-order epistemic reasons are coherent, reasons-responsiveness and Inter-Level Coherence do not lead to incompatible verdicts.

3.2.1. Level-Splitting and Incommensurability

I see two possible explanations of why, in some situations, first-order reasons and higher-order reasons come apart. The first explanation is that higher-order reasons are of a special kind and cannot be compared to first-order reasons. Let's call this the argument from incommensurability, as in the following:

Incommensurability. Epistemic reasons to believe P and epistemic reasons concerning what one has sufficient reason to believe are incommensurable. The balance of epistemic reasons to believe P can differ from the balance of reasons for believing that one has sufficient epistemic reason to believe P.

Here is another way to put it. Let's suppose that first-order reasons are always commensurable with higher-order reasons. In view of the foregoing, reasons to believe that there are reasons to believe P are reasons for believing P, and reasons for believing P are reasons to believe that there are reasons to believe P. So, in a case like Bad Reasoning, Watson should not judge that he has two distinct sets of epistemic reasons (one set of epistemic reasons concerning P and one set of epistemic reasons concerning whether it is rational to conclude that P). He should consider that Holmes's claim that he made a mistake in processing his epistemic reasons is a new reason affecting (to a certain degree) his conclusion that Jack the Ripper is guilty.[88] But

[88]There is ample debate on how much weight Watson should give to Holmes's testimony. This issue is related to recent works on conciliationism in cases of peer disagreement. For arguments in favour of conciliationism, see Christensen (2014) and Feldman (2005b). For arguments in favour of the steadfast view, see Kelly (2005) and

now, suppose that Holmes's testimony is not a reason against the conclusion that Jack the Ripper is guilty, but only a reason to believe that such a conclusion is not supported by epistemic reasons.[89] In such a case, sufficient epistemic reasons could lead to level-splitting. Thus if Incommensurability is true, we would learn something from cases like Bad Reasoning. Indeed, from Watson's perspective, Holmes's testimony could be sufficient evidence to draw a higher-order conclusion, while the various pieces of evidence he gathered could lead him to conclude that Jack the Ripper is the killer. Each type of epistemic reasons could play distinct roles.

Following many others, I find the Incommensurability argument highly implausible.[90] Indeed, suppose that there are cases where higher-order epistemic reasons are not commensurable with reasons for believing P or against believing P. Now, let's assume that an agent has an infallible reason to believe that he or she has sufficient reason to believe P *and* an infallible reason against believing P. Such a situation would not be impossible, since the incommensurability argument implies that higher-order epistemic reasons and first-order reasons are of different kinds. So, an epistemically rational agent could be perfectly confident that he or she has sufficient epistemic reason to believe P, but also be perfectly confident that P is false. As Horowitz rightly stresses, the agent would conclude that whether P and whether he or she has epistemic reasons to believe P are *entirely separate issues*, which appears nonsensical (Horowitz 2014a, 726). Specifically, it is highly implausible that, in some cases, reasons to believe that there are reasons to believe P does not have even the slightest impact on reasons to believe P.

Someone could still argue that incommensurability holds in cases where an agent is mistaken about what the notion of "sufficient reason" means. For example, imagine that Watson rationally believes that the notion of "sufficient reason" refers to "Holmes's testimony" (say, the Detective's dictionary offers that definition and Watson regularly consults this dictionary). In such a case, Watson could be rational in believing that there is sufficient reason to believe P while disbelieving P. Indeed, Holmes could testify to Watson that P while

Schoenfield (2014). See Christensen (2009) for an overview of the debate.
[89]Coates (2012) endorses such a view.
[90]See notably Horowitz (2014a, sec. 3) and Littlejohn (2015, sec. 5).

Watson has "genuine" epistemic reason against believing P. Relative to his belief that sufficient reasons amount to Holmes's testimony, Watson would then be rational in concluding that "there is sufficient reason to believe P, but P is false".[91]

However, such an objection is inconclusive. Inter-Level Coherence implicitly presupposes that the agent correctly uses to the notion of sufficient reason. Specifically, Inter-Level Coherence requires of a *linguistically competent believer* that, if he believes that he has sufficient reason to believe P, he believes P. In the above situation, Watson's use of the notion of sufficient reason is incorrect relative to our linguistic standards. While Watson is rational in concluding "I don't believe that P, but I have sufficient epistemic reason for P", this is explained by the fact that he does not use the notion of sufficient reason we are interested in.

3.2.2. Level-Splitting and Fallible Reasons

If reasons for believing P and reasons for believing that there are reasons for believing P are commensurable, this means that higher-order reasons can somehow count as first-order reasons. The denial of Incommensurability paves the way for various principles connecting higher-order reasons and first-order reasons. Nevertheless, such principles could be correct while cases of level-splitting are possible.[92] So, there must be another explanation of why first-order reasons and higher-order reasons can come apart.

A second explanation of why there could be cases of level-splitting is that higher-order epistemic reasons are *fallible*. We can imagine how higher-order fallible reasons can open the door to cases of level-splitting, as in the following:

Higher-Order Fallibilism. One can have *fallible* sufficient reason for believing that one has sufficient reason to believe P. In a case where such a reason is misleading, it is possible that one is rational to conclude that he or she has sufficient epistemic reason to believe P while lacking sufficient reason for the belief that P.

[91]Littlejohn (2015, 7-8) seems to think that perspectivism about epistemic reasons entails such a problematic conclusion.

[92]As I will explain in section 3.3.2 and 3.4.2, Lasonen-Aarnio (2015, 169) argues that the Rational Reflection principle, which roughly states that an agent's rational expectations of the rational credence in P constrains his or her rational credence in P, can lead to rational epistemic akrasia (see also Elga (2013) on the Rational Reflection principle). However, this principle presupposes that whether P and whether there are epistemic reasons to believe P are not separate issues. So, even if we admit that higher-order reasons and first-order reasons are commensurable, this doesn't seem sufficient to rule out the possibility of level-splitting.

Suppose that an agent has fallible reasons for believing that he or she has sufficient reason to believe P. If higher-order reasons are fallible, having sufficient reason for believing that one has sufficient reason to believe P does not entail the conclusion that one has sufficient reason to believe P, since these reasons could be misleading. So, it is possible that fallible higher-order reasons lead to level-splitting. It seems that Rational Puzzle could be explained by Higher-Order Fallibilism, since one could then be rational to believe that he or she has sufficient epistemic reason to believe P while not believing P (either by withholding judgment on whether P or by disbelieving P).

If Higher-Order Fallibilism is true, we will learn something from cases like Fallible Reasons. Suppose that Watson believes that he has sufficient reason to conclude that Jack the Ripper is the killer. Watson's belief can be based on sufficient epistemic reasons, but not necessarily on infallible epistemic reasons. While such a belief can be rational, it could be based on fallible and misleading reasons. This means that Watson could lack sufficient reasons to draw the conclusion that Jack the Ripper is the killer. Watson would be rational not to conclude that Jack the Ripper is the killer.

It seems that, apart from Incommensurability and Higher-Order Fallibilism, there is no third possible explanation of why Rational Puzzle holds. Indeed, if higher-order sufficient reasons are infallible, having sufficient reason for believing that one has sufficient reason to believe P means that one *inevitably has* sufficient reason to believe P, and so there cannot be cases of level-splitting. Consequently, if Rational Puzzle holds, the culprit is Higher-Order Fallibilism.

3.3. Higher-Order Fallibilism

In this section, I suggest that Rational Puzzle is closely related to other well-known issues concerning fallible reasons. We cannot give a definitive answer to Rational Puzzle without solving traditional problems of responding to fallible reasons, such as the lottery paradox.

3.3.1. Canonical Problems Related to Responding to Fallible Reasons

It is very intuitive to think that rational reasoning has some 'logical properties', such that if you reason correctly from rational beliefs, your conclusion should also be rational. For instance, if the agent is rational in believing P and rational in believing that P implies Q, then he or she should be rational in believing P. After all, the conclusion of a deductive inference cannot be false if its premises are true. As with the preservation of truth from the premises to the conclusion, it could be argued that the conclusion of a deductive inference cannot be irrational if its premises are rational.

This idea underlies many arguments one can find in the literature. For instance, Broome thinks that the following requirement of rationality is correct: "Rationality requires of N that, if N believes at t that p, and N believes at t that if p then q, and if N cares at t whether q, then N believes at t that q" (Broome 2013, 157). Presumably, an implication of the above principle is that, if one is rational in believing P and in believing (P entails Q), it can't be irrational for one to believe Q. In accordance with Broome's requirement, Lord thinks that decisive reasons transmit through deductively valid inference. He suggests that "if the reasons you possess decisively support p and the reasons you possess decisively support if p then q, then the reasons you possess decisively support q." (Lord 2014, 165). Way and Whiting go a step further and argue that it is *constitutive* of correct reasoning to help us reach new rational conclusions. Indeed, they suggest that "if we begin from justified attitudes, we should expect correct reasoning to lead us to further justified attitudes" (Way and Whiting 2016b, 1877).

There are numerous ways to account for such an intuitive result. The most common is to endorse the *rational closure* principles, as in the following:

Rational Closure (under conjunction). Necessarily, if (i) A is epistemically rational in believing P, (ii) A is epistemically rational in believing Q, and (iii) A comes to believe (P^Q) through a good reasoning process (such as deductive reasoning from P and Q), then A is epistemically rational in believing (P^Q).

Rational Closure (under entailment). Necessarily, if (i) A is epistemically rational in believing P, (ii) A is epistemically rational in believing (P entails Q), and (iii) A comes to believe Q through a good reasoning process (such as deductive reasoning from P and P entails Q), then A is epistemically rational in believing Q.

Rational Closure principles are intuitively correct and many philosophers endorse them. However, when combined with Rational Closure principles, the possibility of responding correctly to fallible reasons leads to well-known paradoxes. Here is why. It seems perfectly plausible that rational beliefs are sometimes false (Greco 2014, 203). Specifically, it seems that an agent can be rational in believing P when P's epistemic probability is smaller than 1. For example, if one is certain that P has 0.95 chance (or any other high but imperfect threshold) of obtaining, then one is rationally permitted to believe P. However, the Rational Closure principles seem incompatible with such a verdict, as in the following examples:

Lottery. Imagine a lottery with a sufficiently high number of tickets. Only one ticket is a winner. Each ticket is equally likely to win. Since the probability that each ticket will lose is more than 0.95 (or any other probability that you like), an agent should rationally believe that each ticket is a loser. Indeed, the agent's beliefs concerning chances of winning reflect his or her knowledge of the objective probabilities.[93] However, it is rational to believe that one ticket will win. So, one should believe that each ticket is a loser and that one ticket is a winner, which is inconsistent.

Cheap Justification. Imagine that the sufficient threshold for believing any proposition is 0.95. An agent rationally believes that there is a 0.96 chance that there is a 0.96 chance that P (and a 0.04 chance that there is 0 chance that P). Indeed, the agent's rational beliefs concerning chances reflect his or her knowledge of the objective probabilities.[94] Since the sufficient threshold for believing a proposition is 0.95, the agent then comes to the conclusion that there is a 0.96 chance that P (since, from the agent's perspective, such a proposition has 0.96 chance of obtaining). The agent then comes to the conclusion that P, since (again) the sufficient threshold for believing a proposition is 0.95. However, since 0.92 is equivalent to $\approx 0.96 \cdot 0.96$, the agent is irrational in believing P (since $0.92 \leq 0.95$).[95] So, one is rationally prohibited from believing that P, but can still manage to identify a justificatory chain to the conclusion that P, which is nonsensical.

Various solutions to cases like Lottery and Cheap Justification have been suggested. A first solution is to argue that sufficient reasons are infallible (or may not saliently appear

[93]I here assume that something like the Principal Principle is correct.

[94]See the previous footnote.

[95]At least *in some situations*, such an equivalence is correct. Imagine that an agent is about to roll two dice and that there is 0.92 chance that he or she will not roll a six twice. However, he or she could consider that there are two probabilities here (one for the first die and one for the second). The agent could believe that there is a 0.96 chance that there is a 0.96 chance that he or she will not roll a six twice. Formally, there are different ways to understand this equivalence, but here is a straightforward one. Since $P(B) \cdot P(C|B)$ amounts to $P(B \wedge C)$, it suffices to say that $A=(B \wedge C)$ for it to be rationally permitted to replace $P(A)$ with $P(B) \cdot P(C|B)$. For example, if $P(B) \cdot P(C|B)=0.92$, $P(B) \approx 0.96$, and $A=(B \wedge C)$, then it is correct to conclude that $P(A) \approx 0.96 \cdot 0.96$. See also Worsnip (m.s., sec. 2) on a similar problem.

fallible).[96] An agent can rationally believe that P only if, relative to his or her epistemic reasons, P could not be false. In cases like Lottery, such a solution prohibits a rational agent from believing that each ticket is a loser, since it is possible that one ticket is a winner. In cases like Cheap Justification, if P is uncertain, no infallible justificatory chain leading to the conclusion that P can be identified, since some "residual" uncertainty will remain in any justificatory chain.

Another solution is to abandon Rational Closure principles.[97] While P and Q logically imply (P^Q), rationally believing that P and rationally believing that Q are not necessarily sufficient for rationally concluding that (P^Q). In cases like Lottery, this solution implies that, while I rationally believe that ticket 1 is a loser, that ticket 2 is a loser and so forth, I am not rationally permitted to believe that (ticket 1 is a loser *and* ticket 2 is a loser *and*... ticket n is a loser). In fact, this solution to Lottery entails the denial of Intra-Level Coherence, which roughly states that if an epistemically rational agent believes that P and believes that Q, it is false that he or she believes that ~(P^Q). In cases like Cheap Justification, I may rationally believe that there is a sufficiently high chance that P, but that does not necessarily entail the

[96]See Littlejohn, who argues that there are no justified false beliefs (Littlejohn 2012, 99–102, 121–27). It should be noted that this solution does not exclude degrees of beliefs. Probabilism, for example, is compatible with this view. Under some interpretations of probabilism, a credence is just a percentage of certainty (Sturgeon 2008, 162, n.1). Also, the saliency condition can be interpreted in different ways. Clarke (2013) argues that, while rationally believing P is having a rational credence of 1 in P, rational credences are determined by alternative possibilities one entertains. Leitgeb (2014b) defends the claim that an agent's rational credence in P and the partitioning of possibilities he or she entertains determine the sufficient threshold for believing P. In a lottery case where an agent has rational attitudes concerning *every* ticket, this solution amounts to fixing the sufficient threshold for believing that "ticket n will lose" at 1. Elsewhere, Leitgeb argues that those who deny closure principles can't make sense of the following intuitive principle: "if the agent already believes X, then updating on the piece of evidence X does not change her system of (all-or-nothing) beliefs at all" (Leitgeb 2014a, 783). Indeed, if closure principles are false, there are cases where a rational agent who already believes A and receives evidence for the belief that A changes his or her doxastic attitude towards (A^B), which contradicts the above principle.

[97]Demey (2013), Foley (2009) and Sturgeon (2008) reject closure under conjunction and argue that while agents can rationally believe P and rationally believe Q, it can be rational for them to withhold judgment or disbelieve (P^Q). Kroedel (2011; 2013a; 2013b; 2017) argues that epistemic justification has to do with permissibility, and that since permissions do not agglomerate (being permitted to drink and being permitted to drive does not imply that one is permitted to drink *and* drive simultaneously), rationally believing P and rationally believing Q do not agglomerate and warrant the rational conclusion that (P^Q). See Kiesewetter (2018) for a reply. Relatedly, Easwaran and Fitelson (2015) argue that, from an accuracy-centered perspective, it can be rational to believe P and to believe Q, but to disbelieve (P^Q). Specifically, believing P, believing Q and disbelieving (P^Q) can maximize expected accuracy.

rational conclusion that P, since my belief that there is a sufficiently high chance that P is based on fallible reasons. Indeed, consider the following reasoning:

P1. There is a sufficiently high chance that P.
P2. If there is a sufficiently high chance that P, then P.
C: So, P (deductive inference from P1 and P2).

The conclusion (P) is entailed by P1 and P2. But assuming that Rational Closure is false, an agent who rationally believe P1 and rationally believe P2 might be rational not to believe P. Specifically, if P1 and P2 are supported by fallible reasons, then believing P could be irrational.

3.3.2. Rational Puzzle and Fallible Reasons

The above analysis of fallible reasons sheds light on the conflict between reasons-responsiveness and Inter-Level Coherence. Let's assume for a moment that the first solution to Lottery and Cheap Justification is correct and that sufficient reasons are infallible. This would solve the puzzle of this chapter, since rational agents would be required to respond only to *infallible* reasons. Having *sufficient* reason to believe that one has *sufficient* reason to believe P would amount to having *infallible* reason to believe that one has *infallible* reason to believe P, which would necessarily secure the rational conclusion that P. Thus, there could never be sufficient reason to believe that one has sufficient reason to believe P without there being sufficient reason to believe P.

Now, let's assume that the second solution to Lottery and Cheap Justification is correct, so that rational beliefs do not necessarily ground rational reasoning. In such a context, Inter-Level Coherence would not be a genuine rationality requirement, since one can be rational in believing that one has sufficient reason to believe P, while not believing that P. Consider cases like Cheap Justification. One is rational in believing that there is a 0.96 chance that P. A 0.95 chance that P would constitute a sufficient reason to believe P. Nevertheless, it would be irrational for him or her to believe P, since relative to that agent's epistemic reasons, P has a 0.92 chance of being the case.

Here is another way to put it. Suppose that one is epistemically rational in believing that one has sufficient epistemic reason to believe that P. Suppose furthermore that one is rational in believing the following conditional: if one has sufficient epistemic reason to believe that P, then P. So, the content of one's beliefs ("one has sufficient epistemic reason to believe that P" and "if one has sufficient epistemic to believe that P, then P") entails that P by deductive reasoning. In such a case, one is rationally permitted to conclude that P insofar as Rational Closure under entailment is true. Indeed, if Rational Closure under entailment is true, the conclusion of a deductive inference cannot be irrational if its premises are rational. However, Rational Closure under entailment conflicts with the possibility of responding correctly to sufficient but fallible reasons. Again, assume that a 0.95 chance that P would constitute a sufficient reason to believe P. So, even if there is a 0.96 chance that "one has sufficient epistemic reason to believe that P" and there is a 0.96 chance that "if one has sufficient epistemic to believe that P, then P", this does not entail that P's probability is greater than 0.95. Cases like Cheap Justification support such a conclusion.

Interestingly, some of Lasonen-Aarnio's examples in favour of the conflict between an agent's rational expectations of the rational credence in P and enkratic requirements are very close to cases like Cheap Justification, as she indicates in the following:

> Assume that the threshold for belief is 0.9, and that you know this. Assume that you have the following rational credences: your credence that the rational credence in p is 0.89 is 0.9, and your credence that the rational credence in p is 0.99 is 0.1. Then, your expectation of the rational credence is 0.9.... Given the 0.9 threshold for belief, you believe p. But you also believe that it is not rational to believe p. Hence, you are in a state of epistemic akrasia (Lasonen-Aarnio 2015, 169).[98]

Offering a full solution to Rational Puzzle boils down to determining the constraints on responding to fallible reasons. Rather than being a brand-new puzzle, Rational Puzzle seems

[98]Elsewhere, she offers another example close to Cheap Justification: "Assume, for instance, that *p* is sufficiently likely, and it is only likely to degree 0.3 that *p* is *not* sufficiently likely (and hence, likely to degree 0.7 that *p is* sufficiently likely). Nevertheless, one has misleading evidence about how likely it is that *p* is not sufficiently likely: in fact, it is very likely (say to degree 0.95) that it is likely that *p* is not sufficiently likely.... For all that has been said, the belief *that she is not rationally permitted to believe p* can satisfy the entirety of the above condition" (Lasonen-Aarnio m.s., 5). The difference between her example and mine is this: she identifies sufficient epistemic reason to believe P that are not sufficient epistemic reason to believe that there is sufficient epistemic reason to believe P.

to be a consequence of latent issues concerning fallible reasons. If sufficient reasons are infallible, then there cannot be a dilemma between Inter-Level Coherence and reasons-responsiveness. But if rational beliefs do not necessarily ground rational reasoning, then Inter-Level Coherence could not be a genuine rationality requirement. Thus, as long as we do not have a clear picture of the constraints limiting how agents respond to fallible reasons, we will not be in a position to give a full answer to Rational Puzzle, since Inter-Level Coherence could not be a genuine rationality requirement.

3.3.3. The Objection from the "Same Considerations"

Here is an objection to my argument. I just argued that, under the assumption that sufficient reasons are fallible, a body of epistemic reasons can warrant the second-order belief that one has sufficient epistemic reason to believe P without also warranting the belief that P. But plausibly, the same considerations warrant the belief that P and the belief that one has sufficient epistemic reason to believe P.[99] For example, suppose that I am informed that it is certain that P. This information is both a reason to believe P and a reason to believe that I have sufficient epistemic reason to believe P. In such a context, how can first and higher-order reasons come apart? This would violate the very plausible assumption that the same considerations warrant first and higher-order beliefs. Hence, it seems that there is something wrong with my argument, because we can't imagine a consideration warranting a second-order belief without also warranting a first-order belief.

This objection is tricky. There are two ways to interpret the claim that the same considerations warrant the belief that P and the belief that one has sufficient epistemic reason to believe P. Either (i) the same considerations provide *some* support for the belief that P and the belief that one has sufficient epistemic reason to believe P or (ii) if a consideration provides *sufficient* support for the belief that one has sufficient epistemic reason to believe P, it also provides *sufficient* support for the belief that P. Now, (i) is compatible with the

[99]This objection is inspired by an argument found in Way and Whiting (2016b). According to them, "if you are justified in making no decision regarding φing, it is hard to see how you could also be justified in believing that you ought to φ. Whatever considerations justify you in thinking that you ought to [φ].... would presumably defeat any considerations which might otherwise justify making no decision as to whether to do so" (Way and Whiting 2016b, 1880).

argument of this chapter. Indeed, I do not deny that the same considerations can provide *some* support both at the first and at the higher-order level. My argument is this: under the assumption that sufficient reasons are fallible, the same considerations can *sufficiently* support the belief that one has sufficient epistemic reason to believe P and *insufficiently* support the belief that P. In other words, I merely argue that (ii) is false: even if a consideration provides *sufficient* epistemic reason to believe that one has sufficient epistemic reason to believe P, this consideration does not necessarily provide *sufficient* epistemic reason to believe P.

Cases like Lottery give us a glimpse at what is wrong with (ii). Obviously, the considerations warranting the belief that ticket 1 is a loser and the belief that ticket 2 is a loser have something to do with the considerations justifying the belief that (ticket 1 is a loser *and* ticket 2 is a loser). After all, the content of such beliefs overlap! So, if a consideration X indicates that ticket 1 is a loser and a consideration Y indicates that ticket 2 is a loser, such considerations should have an impact on whether one should believe that (ticket 1 is a loser *and* ticket 2 is a loser).

However, it is possible to have *sufficient* epistemic reason to believe that ticket 1 is a loser and sufficient epistemic reason to believe that ticket 2 is a loser without having *sufficient* epistemic reason to believe that (ticket 1 is a loser *and* ticket 2 is a loser). Here is why. Suppose that I should believe P insofar as P's probability is higher than 0.95. Suppose that the probability that ticket 1 is a loser is 0.96 and that the probability that ticket 2 is a loser is also 0.96. In such a case, one has sufficient epistemic reason to believe that ticket 1 is a loser and sufficient epistemic reason to believe that ticker 2 is a loser, but one lacks sufficient epistemic reason to believe that (ticket 1 is a loser *and* ticket 2 is a loser). Indeed, the probability that (ticket 1 is a loser *and* ticket 2 is a loser) is less than 0.95. In such a case, the same considerations lead an epistemically rational agent to believe that ticket 1 is a loser, that ticket 2 is a loser, but not to believe (ticket 1 is a loser *and* ticket 2 is a loser). Indeed, having sufficient epistemic reason to believe the first two propositions doesn't entail that one has sufficient epistemic reason to believe the conjunct of both propositions.

Cases like Lottery show that *the exact same considerations* can sufficiently support the belief that A, sufficiently support the belief that B but not sufficiently support the belief that (A^B). This paves the way for showing that there is something wrong with the claim that, if a consideration provide *sufficient* support for the belief that one has sufficient epistemic reason to believe P, it also provides sufficient support for the belief that P. Indeed, consider a case like Cheap Justification. Let's assume that an agent knows that the probability that (the probability that P is 0.96) is 0.96. Let's also assume that one has sufficient epistemic reason to believe P insofar as P's probability is higher than 0.95. In such a case, one has sufficient epistemic reason to believe that there is sufficient epistemic reason to believe that P (or to believe that P's probability is greater than 0.95, which is the sufficiency threshold). Yet, this does not exclude that one lacks sufficient epistemic reason to believe P (or that P's probability is less than 0.95). In fact, in Cheap Justification, the *same considerations* explain why (i) the probability that (the probability that P is 0.96) is higher than 0.95 and (ii) P's probability is less than 0.95. Hence, under the assumption that there can be sufficient but fallible reason to believe P, there are counterexamples to (ii).

In summary, it is false to assume that, if a consideration provides *sufficient* epistemic reason to believe that one has sufficient epistemic reason to believe P, this consideration also provides *sufficient* epistemic reason to believe P. This is why the objection from the same considerations fails. Once again, rather than being a brand-new puzzle, Rational Puzzle seems to be a consequence of latent issues concerning fallible reasons.

3.3.4. Implications of the Argument

Following the essential similarities between Cheap Justification and Lottery, we cannot solve the problem of inter-level incoherence without also solving the problem of intra-level incoherence. Accordingly, some well-known vindications of Inter-Level Coherence appear incomplete. Many arguments against rational akrasia stress how *incoherent* it would be to hold an akratic combination of attitudes. Superlatives don't lack to express such an incoherence: akratic combinations of beliefs appear Moore-paradoxical, can lead to patently bad reasoning,

lead to inner conflicts or failures, echo an opaque mindset and so forth.[100] No one will deny that akrasia is a clear type of inter-level incoherence. But we can't just assume that, since akrasia is an incoherent state, then it is an irrational state. One thing I hope to have shown in this chapter is that the possibility of rational incoherence is a latent epistemological issue, especially in puzzles like the lottery paradox. Thus, assuming that incoherent combinations of attitudes are necessarily irrational presupposes a specific solution to these puzzles.

Here is another way to put it. If we are to argue that akrasia is irrational because akratic agents are incoherent, we presuppose that incoherence is necessarily irrational. Now, take a lottery case. Suppose that I believe that ticket 1 is a loser, that ticket 2 is a loser and so forth. What if I refuse to believe that (ticket 1 is a loser *and* ticket 2 is a loser *and*... ticket n is a loser)? Surely, I display some kind of incoherence. I believe that each ticket is a loser, but I refuse to believe the conjunction of all my conclusions concerning the tickets. If one thinks that it is possible to rationally maintain such a state of inner incoherence, then one cannot argue that epistemic akrasia is irrational *on the sole basis* that such a combination of attitudes is incoherent.

Some types of *practical* infallibilism are also affected by the argument of this chapter. Here is a good example. In a recent chapter, Way and Whiting (2016b) argue that if an epistemically rational agent believes that he or she ought to φ, then he or she ought to φ (their argument concerns an epistemically justified agent, but I will assume that justification and rationality are coextensive). This is what they call "Ought Infallibilism". Their argument relies in part on requirements of epistemic rationality. Indeed, much of their argument is related to rational reasoning and the transmission of justification, as we can see with the following step of their argument:

> **(3)** If you reason correctly from justified premise-attitudes, you will reach a justified conclusion-attitude (p. 1876).

Step (3) plays an important role in Way and Whiting's argument. It allows them to conclude that, if an agent justifiably believes that he ought to φ, he has justification for deciding to φ and lacks justification for deciding not to φ (Way and Whiting 2016b, 1880).

[100]See note 78.

But what if reasoning correctly from rational attitudes does not necessarily imply that the conclusion one reaches is rational?[101] As I previously pointed out, one way to solve the lottery paradox is to argue that justified attitudes do not necessarily ground rational reasoning. Consequently, if the right solution to puzzles of fallible reasons is to deny that justified attitudes necessarily ground rational reasoning, Way and Whiting's argument is jeopardized, since an agent could have sufficient reason for believing he ought to φ while lacking sufficient reason for deciding to φ.[102]

Also, the argument of this chapter has implications in the debate surrounding the normativity of epistemic rationality. Indeed, explaining the normativity of requirements such as Inter-Level Coherence has proved to be a difficult task. In response to this challenge, Kiesewetter argues that (i) responding correctly to reasons agents have is normative (Kiesewetter 2017, chap. 7) and that (ii) akratic agents necessarily failed to respond correctly to their epistemic reasons (Kiesewetter 2017, 248–54). This provides an elegant solution to the problem of the normativity of requirements like Inter-Level Coherence: akratic agents are irrational, but Inter-Level Coherence can be derived from a failure to respond correctly to epistemic reasons one has.[103]

However, Kiesewetter acknowledges that intra-level incoherence (as in the Lottery paradox) is not necessarily reducible to a failure to respond correctly to epistemic reasons agents have (Kiesewetter 2017, 254–58). This leads him to argue in favour of a weaker intra-level coherence requirement that is compatible with the fallibilist solution to Lottery, as in the following:

Weak Modus Ponens Coherence. "[Under the assumption that A has sufficient epistemic reason to believe p insofar as A's evidential probability is higher than 1-2ε,] if A believes that p, and A believes that p→q, and A's credence C(p)>1-ε, and A's credence C(p→q)>1-

[101]See also Worsnip (m.s., sec. 2) on a similar worry.
[102]There might be other ways to save the argument in favour of practical infallibilism—perhaps step (3) could be reformulated. My point here is that, so stated, step (3) is contentious. Still, Way and Whiting could identify an alternative route to their conclusion.
[103]To be specific: Kiesewetter (2017, 254) mentions that there could be exceptional cases where one's epistemic reasons support akratic combinations of beliefs. According to him, if such cases were possible, they would show that Reasons Enkrasia is not a genuine requirement of epistemic rationality.

ε, and A attend to q, and A does not believe that q, then A is irrational" (Kiesewetter 2017, 258).

Now, if the argument of this chapter is correct, one can't make a distinction between the problem of intra-level incoherence and the problem of inter-level incoherence. That is, one can't argue on the one hand that epistemic akrasia is necessarily a failure to respond correctly to epistemic reasons one has, but argue on the other hand that some types of intra-level incoherence are not a failure to respond correctly to epistemic reasons one has. Accordingly, Kiesewetter's explanation of the normativity of requirements such as Inter-Level Coherence is compromised, since akratic combinations of beliefs can be supported by an agent's sufficient but fallible epistemic reasons. Hence, since Kiesewetter thinks that there are exceptions to Intra-Level Coherence, he should also accept that there are exceptions to Inter-Level Coherence.

3.4. The Possibility of Always Responding to Higher-Order Infallible Reasons

I just argued that we cannot offer a full solution to Rational Puzzle without also solving canonical problems of responding to first-order fallible reasons. Nevertheless, I do believe that there is one important difference between responding to first-order fallible reasons and responding to higher-order fallible reasons—that is, under one interpretation of higher-order reasons, there is no obstacle to eliminating higher-order fallible reasons. My argument relies on the probabilistic representation of reasons introduced in section 1.1 and can be roughly summarized as follows:

(1) There can be cases of level-splitting only if agents respond to higher-order fallible reasons.

(2) Relative to the probabilistic representation of reasons, the weight of higher-order fallible reasons can be represented as conditional probabilities and the weight of higher-order infallible reasons can be represented as unconditional probabilities.

(3) So, there can be cases of level-splitting only if the weight of reasons agents have is represented as conditional probabilities.

(4) But conditional probabilities can be replaced by unconditional probabilities.

(5) So, relative to the probabilistic representation of reasons, fallible higher-order reasons can be replaced by infallible higher-order reasons, and agents can avoid responding to fallible higher-order reasons.

(C) So, relative to the probabilistic representation of reasons, cases of level-splitting can be avoided.

Consequently, there is no reason why a rational agent would necessarily have to choose between being reasons-responsive and satisfying Inter-Level Coherence. Furthermore, it seems plausible that a rational agent would prefer to ground his or her beliefs concerning what he or she has sufficient reason to believe on infallible reasons. In summary, I do not rule out the possibility that a rational agent can maintain an akratic combination of attitudes while responding correctly to his or her epistemic reasons, but I claim that this would be an odd preference.

3.4.1. Conditional and Unconditional Probabilities

Let's assume for a moment that the rational status of Inter-Level Coherence is uncertain and that we cannot give a full answer to Rational Puzzle. In view of the foregoing, what are we in a position to defend? I previously argued that if all higher-order reasons are infallible, then there cannot be cases of level-splitting. This means that there are two ways to offer a partial solution to Rational Puzzle, as in the following:

(1) While there are first-order fallible reasons, higher-order reasons concerning facts about reasons or rationality are infallible.[104]

(2) While it is possible for an epistemically rational agent to respond to higher-order *fallible* reasons, he or she is always in a position to respond to higher-order *infallible* reasons.

I will now provide an argument for (2), the claim that an agent is always in a position to respond to higher-order infallible reasons. This provides a partial solution to Rational Puzzle,

[104]This view is very close to Titelbaum's (2015) Fixed Point thesis, which roughly states that mistakes concerning the requirements of rationality are mistakes of rationality. However, Titelbaum's Fixed Point thesis relies on the premise that akrasia is irrational (Titelbaum 2015, 254), an assumption that I question in this chapter. Also, the claim that mistakes concerning the requirements of rationality are mistakes of rationality is compatible with the rejection of reasons-responsiveness. Consider the following argument: (1) Rational agents cannot be mistaken concerning what rationality requires of them; (2) however, in responding correctly to their reasons, agents can form rational false beliefs concerning what they sufficient reason to believe; (C) so, responding correctly to reasons an agent has is not a genuine requirement of rationality, or claims concerning reasons-responsiveness are outside the realm of rationality. For these reasons, I will not explore Titelbaum's line of reasoning here. However, I acknowledge that exploring such a line of reasoning could eventually solve Rational Puzzle.

since if one can avoid responding to higher-order fallible reasons, then one is always in a position to satisfy both Inter-Level Coherence and Reasons-Responsiveness.

I previously assumed that epistemic reasons warrant *epistemic probabilities*, understood as the probabilities warranted by an agent's body of epistemic reasons. With respect to Rational Puzzle, we can learn something from such a representation of reasons.

There are two main types of probability assessments—namely, conditional probabilities and unconditional probabilities. In other words, we can wonder what P's unconditional probability is, but we can also wonder what P's probability is on the condition that some states of affair (Q, R, S...) obtain. A quick clarificatory remark: we could also say that an unconditional probability is a probability conditional on a necessarily true event or proposition. For example, since Qv~Q is necessarily true, then Pr(P)=Pr(P|(Qv~Q). In other words, P's unconditional probability can be understood as P's probability conditional on a proposition with probability 1 (like Qv~Q). In the remainder of this chapter, I will use this technical sense of "unconditional probability."

Suppose some first-order proposition P is supported by fallible reasons. Given the probabilistic representation of the weight of reasons, this means that Pr(P)=X (where X<1). Since P's probability is warranted by a body of epistemic reasons, we can "conditionalize" P's probability on the reasons one has. Let Q denotes such epistemic reasons. Then, we get that Pr(P|Q)=X. Q can, in turn, be supported by infallible or fallible higher-order epistemic reasons. If Q is supported by infallible epistemic reasons, P's probability is technically unconditional—that is, if Q is supported by infallible epistemic reasons, Pr(P|Q)=Pr(P|Qv~Q). However, if Q is not supported by infallible epistemic reasons, P's probability is conditional on the condition that Q obtains, and Q's probability is not 1.[105]

[105]Alternatively, Q could be certain, but the support relation Pr(P|Q)=X could be uncertain (in other words, it could be uncertain that P's probability on Q is X). Then, Pr(Pr(P|Q)=X)<1. This amounts to saying that, if some condition R obtains, Pr(P|Q)=Y, where X≠Y. Then, we can simply conditionalize the support relation on R, as in: Pr(P|Q^R)=Y and Pr(P|Q^~R)=X. As we can see, framing the problem in terms of uncertainty about the support relation turns out to be notationally equivalent to uncertainty concerning the possession of some reasons. It's a different way to interpret the notation.

What the above remarks suggest is that the weight of fallible higher-order reasons with conditional epistemic probabilities. If the probability that [P's probability is 0.9] is 0.9, then P's probability is 0.9 on the condition that Q obtains, and Q's probability is 0.9. In such a case, it could be false that P's probability is 0.9, since such a claim is conditional on Q obtaining, and Q is uncertain. By way of contrast, the weight of infallible higher-order reasons can be represented by unconditional epistemic probabilities. If it is certain that P's probability is 0.9, then such an evaluation of P's probability is not conditional on some merely probable event Q obtaining (it can be conditional on some event Q obtaining, but Q's probability has to be 1).

One reason why it is appropriate to represent higher-order reasons by conditional and unconditional epistemic probabilities is that such a representation is compatible with the Commensurability constraint discussed in section 3.2 (according to such a constraint, higher-order reasons can count as first-order reasons). Here is why. Suppose that P's epistemic probability is 0.9 on the condition that Q obtains and that P's epistemic probability is 0 on the condition that ~Q obtains. P's probability will vary depending on Q's obtaining. In particular, if Q were certain, this would entail that P's probability is 0.9. Similarly, if ~Q were certain, this would entail that P's probability is 0. As we can see, the existence of reasons for or against the conclusion that Q can affect the probability of first-order conclusions such as P. Since epistemic reasons are represented by epistemic probabilities, we can conclude that acquiring higher-order epistemic reasons can somehow count as acquiring first-order reasons. Hence, the Commensurability condition discussed in section 2 is satisfied.

Here is the trick: as long as chains of conditional probabilities end with an unconditional probability, a conditional probability can be replaced by an unconditional probability. For example, if the epistemic probability that [the epistemic probability that P is 0.9] is 0.9 and the epistemic probability that [the epistemic probability that P is 0] is 0.1, it is possible to determine P's unconditional epistemic probability. For example, in this specific case, P's unconditional epistemic probability would be 0.81.[106] In other words, the epistemic

[106]We can express such a result formally. Suppose that, conditional on A, P's probability is X, but conditional on ~A, P's probability is Y. Conditions A and ~A are also merely probable. Let's assume that Pr(P|A)=X, Pr(P|~A)=Y, Pr(A)=C and Pr(~A)=(1-C). In such a context, we can determine P's conditional probability, but we can also determine P's *unconditional* probability. Indeed, Pr(J)=Pr(J^K)+Pr(J^~K) and Pr(J^K)=Pr(K)·Pr(J|K) are

probability that [the epistemic probability that P is 0.81] is 1. Now, recall that the weight of infallible higher-order reasons can be represented by unconditional epistemic probabilities. This means that, all things being equal, we can pass from higher-order fallible reasons (as represented by conditional epistemic probabilities) to higher-order infallible reasons (as represented by unconditional epistemic probabilities). That is, the same body of epistemic reasons can be understood as providing higher-order *fallible* reasons and higher-order *infallible* reasons.

We can move from conditional epistemic probabilities to unconditional epistemic probabilities as long as chains of conditional probabilities end with an unconditional probability. What about the cases where P's epistemic probability is infinitely conditional? For example, there could be cases where P's probability is conditional on Q, and Q's probability is conditional on R, and such a regress does not stop with a "final" unconditional probability. Even in such situations, there is a *modest* sense in which we can move from higher-order fallible reasons to higher-order infallible reasons. Indeed, imagine that P's probability is determined by the following series:[107]

Suppose $Pr(P)=Pr(A_1)-Pr(A_2)-Pr(A_3)...-Pr(A_n)$, where $Pr(A_n)=0.9\cdot(10^{1-n})$.

If $Pr(P)=0.9-\sum_{n=1}^{\infty} 0.9\cdot(10^{-n})$, Pr(P) converges to 0.8.

As we can see, P's probability is here defined by an infinite series of merely probable events, but still converges to 0.8. The lesson here is that while P's probability is conditional on a series of merely probable events, there is a modest sense in which we can determine P's unconditional probability, since P's unconditional probability *converges to* 0.8. If such an infinite probabilistic chain converges, then there is a *modest* sense in which P's unconditional probability can be determined.[108]

familiar probability rules. Since Pr(J^K)=Pr(K)·Pr(J|K), we can conclude that X·C=Pr(P^A) and Y·(1-C)=Pr(P^~A). Since Pr(J)=Pr(J^K)+Pr(J^~K), we can conclude that Pr(P)=(Y·(1-C))+(X·C). In the situation described, since Pr(P|A)=0.9, Pr(P|~A)=0, Pr(A)=0.9 and Pr(~A)=0.1, we get the result that Pr(P)=(0·0.1)+(0.9·0.9)=0.81. Hence, at least in the situation described, combinations of conditional probabilities can be replaced by an unconditional one.

[107]This example is largely inspired by Atkinson and Peijnenburg's (2006; 2009) result that an infinite probabilistic chain can ground P's probability.

[108]For the sake of simplicity, I here limit myself to cases where an infinite chain of conditional probabilities is represented by a convergent series, not a divergent one.

This is an important step toward solving Rational Puzzle. Relative to the probabilistic representation of reasons, higher-order fallible reasons can be represented by conditional epistemic probabilities and higher-order infallible reasons can be represented by unconditional epistemic probabilities. Since conditional probabilities can be replaced by unconditional probabilities, fallible higher-order reasons can be replaced by infallible higher-order reasons, and so it is rational for agents to avoid responding to fallible higher-order reasons. There is no specific reason why it would be necessary for agents to respond to higher-order fallible reasons. Furthermore, if agents can avoid responding to higher-order fallible reasons, cases of level-splitting can also be avoided. This provides a partial solution to Rational Puzzle.

3.4.2. A Step Further: The Conflict Between the Rational Reflection Principle and Enkrasia

The argument I just offered can shed light on the putative conflict between the Rational Reflection principle and enkratic requirements. The Rational Reflection principle roughly states that an agent's rational expectations of the rational credence in P constrain his or her rational credence in P. Lasonen-Aarnio (2015, 169) claims that satisfying the Rational Reflection principle can lead to forming akratic combinations of attitudes. This is so because one can rationally believe P while rationally believing that one's own belief is irrational, as in the following line of reasoning:

(1) It is rational for A to believe P if and only if A has a rational credence of at least 0.9 in P.
(2) The rational credence that [the rational credence in P is 0.89] is 0.9, and the rational credence that [the rational credence in P is 0.99] is 0.1.
(3) Following the Rational Reflection principle, $Cr(P)=(0.99\cdot0.1)+(0.89\cdot0.9)=0.9$, and so A rationally believes P.
(4) But the credence in [the rational credence in P is 0.89] is 0.9. So, A rationally believes that the rational credence in P is 0.89 and that believing P is irrational.

However, Lasonen-Aarnio assumes that credence assignments are rational only insofar as they track (or reflect) epistemic probabilities (Lasonen-Aarnio m.s., 2). This means that, in the above situation, the epistemic probability that [P's epistemic probability is 0.89] is 0.9 and the epistemic probability that [P's epistemic probability is 0.99] is 0.1. Now, if the epistemic

probability that [P's epistemic probability is 0.89] is 0.9, this means that P's epistemic probability is 0.89 *conditional on* a state of affairs Q obtaining, and Q's epistemic probability is 0.9. Similarly, if the epistemic probability that [P's epistemic probability is 0.99] is 0.1, this means that P's epistemic probability is 0.99 *conditional on* Q's not obtaining, and ~Q's epistemic probability is 0.1. Finally, we can use P's conditional probabilities to calculate P's unconditional probability. In the above case, P's unconditional epistemic probability is 0.9 (since (0.89·0.9)+(0.99·0.1)=0.9). This means that the epistemic probability that [P's epistemic probability is 0.9] is 1.

Now, recall that the weight of infallible higher-order reasons are represented by unconditional epistemic probabilities. In such a context, since 0.9 is P's unconditional epistemic probability, it would be rational for an agent to be certain that 0.9 is the rational credence in P. In other words, he or she has an infallible reason to conclude that 0.9 is the rational credence in P, and so being certain that 0.9 is the rational credence in P would be an appropriate response to his or her epistemic reasons. There is no need for the agent to believe that such a credence assignment is irrational relative to his or her epistemic reasons. That is, the agent doesn't need to be mistaken about his or her own epistemic rationality.

Here is another way to put it. In the described case, an agent's rational credences can track the following epistemic probabilities: the epistemic probability that [P's epistemic probability is 0.89] is 0.9 and the epistemic probability that [P's epistemic probability is 0.99] is 0.1. As long as sufficient reasons can be fallible, tracking these epistemic probabilities can lead to a conflict between the Rational Reflection principle and enkratic requirements. However, an agent's rational credences can also track the following epistemic probability: the epistemic probability that [P's epistemic probability is 0.9] is 1. If the agent's rational credences track this epistemic probability, we get the following result:

(5) It is rational for A to believe P if and only if A has a rational credence of at least 0.9 in P.
(6) The rational credence that [the rational credence in P is 0.9] is 1.
(7) Following the Rational Reflection principle, Cr(P)=(0.9·1)=0.9, and so A rationally believes P.
(8) Since the credence in [the rational credence in P is 0.9] is 1, A rationally believes that believing P is rational.

As we can see, when tracking higher-order infallible reasons (as represented by unconditional epistemic probabilities), the Rational Reflection principle does not lead to forming an akratic combination of beliefs.

Now, perhaps we should not accept the Rational Reflection principle (Lasonen-Aarnio (2015) ultimately rejects such a principle). I am not defending such a principle here. What I wish to stress is that, when taking the possibility of responding to higher-order infallible reasons into account, the conflict between the Rational Reflection principle and enkratic requirements is a lot less clear. Surely, when agents respond to higher-order fallible reasons, the Rational Reflection principle can conflict with enkratic requirements. However, as long as it is possible for the agent to avoid responding to higher-order fallible reasons (which is always the case), such a conflict is resolved.

3.4.3. The Relevance of Responding to Higher-Order Infallible Reasons

If agents are always in a position to respond to higher-order infallible reasons, this means that, minimally, it is *always possible* to satisfy Inter-Level Coherence and be reasons-responsive. I will now go a step further and suggest that rational agents *prefer* responding to infallible higher-order reasons. While this will *not prove* that Inter-Level Coherence is a genuine rationality requirement, such an argument will make it *plausible* that Inter-Level Coherence is a requirement of rationality, since an agent would have no reason to entertain an epistemically akratic combination of attitudes.

Responding to higher-order infallible reasons provides a better answer to cases like Cheap Justification. Recall that, in Cheap Justification, an agent rationally believes that there is a 0.96 chance that there is a 0.96 chance that P (and a 0.04 chance that there is 0 chance that P), and such rational beliefs concerning chances reflect his or her knowledge of the objective probabilities. If cases like Cheap Justification support incoherentism, it must be admitted that a rational agent can frequently figure out a misleading chain of justification in favour of numerous higher-order beliefs concerning sufficient reasons. For example, in some situations where I know that P's objective probability is 0.75, I could believe that there is a ≈0.87 chance

that there is a ≈0.87 chance that P, since 0.75 is equivalent to ≈0.87·0.87.[109] Assuming that 0.85 is a sufficient probabilistic threshold, I could then come to the conclusion that there is a ≈0.87 chance that P. But there's *something quite wrong* with such a result. I take it as a datum that no rational agent would want to have such a misleading justificatory chain of attitudes concerning sufficient reasons. Plausibly, if I know that P's objective probability is 0.75, I am better off believing that there is a 0.75 chance that P, and this seems best explained by the fact that I should respond to infallible higher-order reasons.

Here is another way to understand my point. Allowing fallible higher-order reasons can lead to patently strange situations that no rational agent would want to be in (especially since they can easily be avoided). Consider the following conversation:

Watson: What are the odds that Jack the Ripper did it?

Holmes: You may rationally believe that there is a ≈0.87 chance that Jack the Ripper did it.

Watson: Why would it be rational for me to believe that there is a ≈0.87 chance that Jack the Ripper is guilty?

Holmes: Well, let's see. Undoubtedly, there is a 0.75 chance that Jack the Ripper did it, but in this specific case there is a ≈0.87 chance that there is a ≈0.87 chance that Jack the Ripper did it (and a 0.13 chance that there is 0 chance that Jack the Ripper did it). The sufficient threshold for rationally believing a proposition is 0.85. So, it is rational for you to conclude that there is a ≈0.87 chance that Jack the Ripper is guilty.

Watson: Okay, and so following the same explanation you just provided, I am also rational in concluding that Jack the Ripper did it.

Holmes: No! Your belief that there is a ≈0.87 chance that Jack the Ripper did it is not a sufficient reason for believing that Jack the Ripper did it. You see, since there is no doubt that the objective probability that Jack the Ripper did it is 0.75, you are not permitted to believe that Jack the Ripper is the killer.

Watson: Oh, so you first gave me information from which I cannot rationally reason, but you had information from which I could reason. You gave me a sufficient reason to believe something from which I could badly reason.

Holmes: Exactly!

What do we learn from the above conversation? Even if we assume that Watson did not violate any rule of rationality in believing that there is a ≈0.87 chance that P, it is patently clear to him that, in believing that there is a 0.75 chance that P, he has access to a more

[109]See note 93.

informative and useful way to reason. Believing that there is a 0.75 chance that P would ground correct reasoning, while believing that there is a ≈0.87 chance that P will not. Also, while Holmes is not making any rational mistake in presenting the chances differently, there is a better way for him to inform Watson of P's likelihood. Thus, in situations where fallible reasons concerning what is probable can be replaced with infallible reasons concerning what is probable, the latter appears preferable.

In summary, since beliefs concerning sufficient reasons often aim at reasoning correctly, a rational agent would prefer responding to infallible reasons concerning what he or she has sufficient reason to believe. Furthermore, there seems to be no structural obstacle to avoid responding to higher-order fallible reasons. So, it is possible that an epistemically akratic combination of attitudes is rational, but the higher-order belief that one has sufficient reason to believe P would play no role in an agent's reasoning (or a potentially misleading role). Even if, strictly speaking, it would not be irrational to respond to fallible higher-order reasons, I see no reason why an agent would prefer responding to fallible higher-order reasons.

3.4.4. Limits of the Partial Solution

Since an agent is always in a position to respond to higher-order infallible reasons, he or she never has to choose between satisfying Inter-Level Coherence and being reasons-responsive. Furthermore, since reasoning from infallible higher-order reasons appears preferable to an epistemically rational reasoner, I see no reason why an agent would reason from fallible higher-order reasons and end up with an akratic combination of attitudes. This is why I *partially* solved Rational Puzzle: Inter-Level Coherence and reasons-responsiveness do not necessarily conflict with each other.

Nevertheless, Inter-Level Coherence could fail to be a genuine rationality requirement, since strictly speaking, I did not prove that inter-level incoherence is necessarily irrational. First, the partial solution I put forth relies on the probabilistic representation of the weight of reasons. I assumed thatwe could reach similarresults through other representations of the weight of epistemic reasons, such as possibility theory or ranking theory. But as I indicated in section 1.5.3, the probabilistic representation of the weight of reasons raises methodological

difficulties. Also, assuming such an equivalence between these various representations of the weight of reasons will be unsatisfactory to many philosophers. To reach a complete solution we should either prove that the results of this chapter can be reached through any representation of reasons or adapt the argument to other frameworks.

Second, we face a methodological problem for proving that Inter-Level Coherence is a genuine requirement of epistemic rationality. The problem is this. On the one hand, it is tempting to think that Intra-Level Coherence is not a genuine requirement of epistemic rationality, because of lottery cases and the like. But lottery beliefs are incoherent. In view of the foregoing, there doesn't seem to be a necessary connection between coherence and epistemic rationality.

Assume there is no necessary connection between coherence and epistemic rationality. Then, how can we argue that Inter-Level Coherence is a genuine requirement of epistemic rationality? The conclusion can't follow from the fact that it's a coherence requirement (this would contradict the rejection of Intra-Level Coherence). So, we need another explanation of why Inter-Level Coherence is a genuine requirement of epistemic rationality.

I have argued that rational agents can avoid akratic combinations of beliefs, and that preferring akratic beliefs over coherent combinations of beliefs would be *odd*. Is that enough to conclude that Inter-Level Coherence is a genuine requirement of epistemic rationality? It doesn't seem so. A necessary connection between oddness and rationality is unheard of.

Here is another way to tackle the problem. The Reliability Criterion, introduced in section 1.5.4, roughly states that ideally rational agents take the available standards that optimize their ratio of true to false beliefs. The Reliability Criterion could explain what's irrational with akratic combinations of beliefs. Consider the following reasoning:

(1) In the right conditions, ideally rational agents take the available standards that optimize their ratio of true to false beliefs (Reliability Criterion).

(2) Rational agents have a choice between responding to infallible higher-order reasons or responding to fallible higher order reasons (as argued in §3.4.1).

> (2.1) All things being equal, if one responds correctly to infallible higher-order reasons, one forms true higher order beliefs.

(2.2) All things being equal, if one responds correctly to fallible higher-order reasons, one can form false higher order beliefs.

(3) All things being equal, forming false beliefs reduces one's reliability level.

(4) Following (1)-(3), ideally rational agents who have a choice between responding to infallible higher-order reasons or responding to fallible higher order reasons choose the former.

(5) One can't violate Inter-Level Coherence if one always responds correctly to infallible higher-order reasons.

(C) So following (4) and (5), ideally rational agents do not violate Inter-Level Coherence.

This argument is plausible. It supports the conclusion that rational agents will satisfy Inter-Level Coherence. However, the argument doesn't prove that Inter-Level Coherence is a genuine requirement of epistemic rationality. The fact that X is satisfied by rational agents does not entail that X is rationally required. Compare: ideally rational agents breathe. But breathing is not a genuine requirement of epistemic rationality. There is an explanatory gap between (i) the fact that Inter-Level Coherence is satisfied by rational agents and (ii) the tentative claim that Inter-Level Coherence is rationally required. My point here is this: it is difficult to prove *out of doubt* that Inter-Level Coherence is a genuine requirement of rationality.

This brings us back to the difficulties presented at the beginning of chapter 1. We lack an effective method for determining the requirements of epistemic rationality. Accordingly, even if akratic combinations of beliefs are odd, and even if we can use the Reliability Criterion to argue that rational agents do not violate Inter-Level Coherence, it is still hard to prove that Inter-Level Coherence is a genuine requirement of epistemic rationality.

3.5. Conclusion: Back to the Modest Reductionist Hypothesis

We have covered a lot of ground in this chapter. I have suggested that Rational Puzzle is essentially related to traditional problems of responding to fallible reasons such as the lottery paradox. However, I haved also suggested that an agent never has to choose between responding to his or her reasons and avoiding akratic combinations of attitudes, because he or she is always in a position to satisfy both.

This chapter answers a worry discussed in section 1.2.6, which has to do with unsolvable normative dilemmas of epistemic rationality. According to such a worry, the conflict of epistemic reasons and coherence puts us in an uncomfortable position. However, following the argument of this chapter, Inter-Level Coherence and reasons-responsiveness do not give rise to unsolvable dilemmas of epistemic rationality. Indeed, an agent is always in a position to satisfy both requirements. So, these requirements do not give rise to an objection against the normativity of epistemic rationality.

More importantly, the argument of this chapter partly confirms the Modest Reductionist Hypothesis. Here is why. Rational Closure (under conjunction) states that, if one rationally believes that P and rationally believes that Q, then one is rationally permitted to believe that (P^Q). Following the argument of this chapter, we don't know if Rational Closure is true of false. All we know is this: either it is true, or it is false.

Suppose that Rational Closure is true. If Rational Closure is true, then violating Inter-Level Coherence or Intra-Level Coherence amounts to having an inconsistent combination of beliefs. So Inter-Level Coherence and Intra-Level Coherence can be derived from Consistency. This was made clear in section 3.3.1. Moreover, there exists an argument connecting Consistency and reasons-responsiveness which relies on Rational Closure. This argument has been defended by Andrew Reisner (2011). It goes as follows.

If Rational Closure is true, Consistency can be derived from responding correctly to a specific type of epistemic reason, namely the reasons we have not to believe *a priori* false propositions. So, under the assumption that Rational Closure is true, inconsistency can be understood as a failure to respond correctly to epistemic reasons. To see why, consider the following line of reasoning:

(1) (P^~P) is *a priori* false. If (P^~P) is *a priori* false, agents have conclusive reason not to believe (P^~P);

(2) if an agent has conclusive epistemic reason not to believe (A^B), he or she has conclusive epistemic reason not to believe A and believe B simultaneously;

(C) following (1) and (2), an agent has conclusive epistemic reason not to believe P and believe ~P simultaneously.

As we can see, (2) presupposes Rational Closure. One way to see this is to consider a lottery case. Indeed, suppose that, relative to a lottery of n tickets, an agent has conclusive epistemic reason to believe that not all tickets are losers. We can express such a requirement by the following:

(3) an agent has conclusive epistemic reason not to believe that (ticket 1 is a loser ^ ticket 2 is a loser ^ ticket 3 is a loser ... ticket n is a loser).

Now, combining (2) and (3) implies the following:

(4) an agent has conclusive epistemic reason not to (believe that ticket 1 is a loser & believe that ticket 2 is a loser & believe that ticket 3 is a loser ... & believe that ticket n is a loser).

However, as it was just made clear, (4) holds only insofar as Rational Closure is true. Following the incoherentist solution to Lottery, an agent should be permitted to believe that ticket 1 is a loser and to believe that ticket 2 is a loser ... and to believe that ticket n is a loser. The agent is simply not permitted to derive from his or her rational beliefs that (ticket 1 is a loser ^ ticket 2 is a loser ^ ticket 3 is a loser ... ticket n is a loser). So, we should accept (3) without accepting (4), which means that (2) is false. Hence, (2) is directly connected to Rational Closure. Reisner's argument works only insofar as Rational Closure is true.

Now, suppose that Rational Closure is true, and so that the above line of reasoning is correct. Accordingly, ending up with an inconsistent combination of attitudes amounts to not responding correctly to epistemic reasons agents have. Inconsistent agents would fail to respond to their reasons not to believe *a priori* falsehoods. To make a long story short, if Rational Closure is true, Consistency (and thus Inter-Level and Intra-Level Coherence) can be derived from responding correctly to a specific type of epistemic reasons, namely, the reasons we have not to believe *a priori* false propositions.

Now, suppose that the Rational Closure is false. As I explained at length in section 3.3, if Rational Closure is false, Inter-Level Coherence and Intra-Level Coherence are not genuine

rationality requirements, and so there is no need to explain such requirements in terms of reasons-responsiveness.

We don't know if Rational Closure is true. But this doesn't matter, because the above remarks provide a partial confirmation of the Modest Reductionist Hypothesis: regardless of whether the Closure Conjecture is correct, Inter-Level Coherence and Intra-Level Coherence play no distinct explanatory role in theories of epistemic rationality. Indeed, if Rational Closure is true, Inter-Level Coherence and Intra-Level Coherence are both derivable from responding correctly to reasons one has. If Rational Closure is false, Inter-Level Coherence and Intra-Level Coherence are not genuine rationality requirements. Either way, the Modest Reductionist Hypothesis is confirmed in part: when compared with substantive requirements of epistemic rationality, putative requirements such as Inter-Level Coherence or Intra-Level Coherence have no distinct explanatory role in the theory of epistemic rationality.

Still, there could be another structural requirement playing an explanatory role in the theory of epistemic rationality, namely, Consistency. This brings us to the next chapter.

Chapter 4. The Explanatory Role and the Normativity of Consistency

Chapter summary. Is inconsistency a mere symptom of having violated other requirements of rationality—notably, reasons-responsiveness requirements? Or is inconsistency irrational on its own? This question has important implications for the debate on the normativity of epistemic rationality. In this chapter, I defend a new account of the explanatory role of the requirement of consistency. Roughly, I will argue that, in cases where an epistemically rational agent is permitted to believe P and also permitted to disbelieve P (relative to a body of epistemic reasons), the consistency requirement plays a distinct explanatory role. I will also argue that such a type of permissiveness is a live possibility when it comes to rational epistemic standards.

Linda believes that P and she also believes that ~P. She displays a kind of epistemic irrationality—she is inconsistent. Here is how we can express the requirement she violates:

Consistency. Epistemic rationality requires that, if A believes that P at time t, A does not believe that ~P at time t.

Perhaps you think that inconsistent agents can be rational.[110] After all, in the previous chapter I have argued that, if Rational Closure is false, some incoherent combinations of beliefs are rational. So, why not think that inconsistency can also be rational? I can't prove that Consistency is a genuine requirement of rationality. As I explained in chapter 1, we lack a good method for identifying the requirements of rationality, including this one. If you think that Consistency is not a genuine requirement of rationality, there is no point in reading the remaining two chapters of this thesis. Without Consistency, rationality consists solely in responding correctly to reasons one has, and it would be wildly implausible that any further structural requirements do any distinct explanatory work in the theory of epistemic rationality.

[110]Broome notes that those who think that some contradictions can be true might want to reject Consistency. He writes: "I would not object to weakening the formulae in a suitable way... I expect a similar case could be made for weakening most of the [putative structural] requirements." (Broome 2013, 155) See also Caie (2013) on the related topic of rational probabilistic incoherence.

But if you believe (as I do) that Consistency is a requirement of epistemic rationality, keep reading.

Assuming that Consistency is a requirement of epistemic rationality, some important questions remain unanswered. For instance, *why* is Linda irrational? Does her violating Consistency *explain* why she is irrational? Perhaps her inconsistency is just a *symptom* or a *by-product* of her having violated other requirements of epistemic rationality. Indeed, there is a common explanation of why Linda is irrational:

Strong Reasons-Responsiveness Thesis. Epistemic rationality consists in responding to epistemic reasons one has. No body of epistemic reasons can simultaneously support the belief that P and the belief that ~P. So, requirements such as Consistency have no distinct explanatory role: Linda is irrational in believing P and believing ~P because she is not responding correctly to her epistemic reasons.[111]

The above response has important implications for the debate surrounding the *normativity* of epistemic rationality. Some authors such as Kolodny (2005; 2007b; 2008b) have suggested that there is no reason to be consistent (at least, offering a convincing argument in favour of the normativity of Consistency has proved to be very difficult). However, if Consistency is a mere by-product of other requirements such as responding correctly to epistemic reasons, there is no need to discuss the normativity of Consistency, since this requirement of epistemic rationality would have no distinct explanatory role.

This chapter argues that the Strong Reasons-Responsiveness Thesis is incorrect. Specifically, I will argue that the Strong Reasons-Responsiveness Thesis has a major blind spot—that is, it assumes that a specific type of permissiveness (or permissivism) concerning epistemic reasons is false. Ultimately, I will argue that Consistency provides an explanation of why agents are irrational in cases where both believing P and disbelieving P are warranted by a body of epistemic reasons. In other words, I will argue for the following view:

[111]This strategy has been pursued by Way (2009), McHugh and Way (2017), Kiesewetter (2017, chap. 7) and Lord (2017). From an "internalist" perspective, such a view amounts to substantive internal coherence requirements between (i) *a priori* knowledge and phenomenal experiences and (ii) beliefs or credences—see Wedgwood (2017, sec. 0.5). See also Schroeder (2008; 2011). In the practical realm, the view that rationality consists in responding to reasons one has is often associated with Raz (1999; 2005).

Modest Reasons-Responsiveness Thesis. Epistemic rationality consists (in part) in responding to epistemic reasons one has. No *impermissive* body of epistemic reasons supports both the belief that P and the belief that ~P. So, Consistency plays no distinct explanatory role *in impermissive situations*. However, in cases where both believing P and disbelieving P are warranted by a body of epistemic reasons, Consistency plays a distinct explanatory role.

In section 4.1, I will offer a brief overview of the debate surrounding the explanatory role of Consistency. In section 4.2, I will explain why the Strong Reasons-Responsiveness Thesis presupposes that agents never find themselves in epistemically permissive situations, which is an important blind spot. I will then argue in favour of the Modest Reasons-Responsiveness Thesis. In section 4.3, I will refine my view in response to Kolodny's decision-theoretic argument.

For the sake of brevity, this chapter will make one assumption. I will focus on cases where P's truth is "mind-independent." Some arguments in favour of permissiveness concern cases where believing P guarantees P's truth (Kopec 2015; Raleigh 2015). Such cases are interesting, but they have little to do with the debate addressed in this thesis.[112] So, I will ignore this possibility here.

4.1. The Debate Surrounding the Explanatory Role of Consistency

4.1.1. Elimination and Consistency

Consistency is a structural requirement of epistemic rationality. By way of contrast, there are putative *substantive* requirements of epistemic rationality, which govern how agents form and revise their beliefs in response to their epistemic reasons.

Perhaps Consistency is just a by-product of these other substantive requirements of rationality. Call this the Elimination thesis, expressed as follows:[113]

[112]Specifically, much of the debate has to do with the fact that Consistency is not "truth-conducive" (see section 1.2.5). Naturally, such an objection is pointless in cases where believing P guarantees that P will be true.

[113]Broome (2005; 2007a; 2007b; 2013, chap. 5) and Worsnip (2015) reject Elimination. Kolodny (2005; 2007a; 2007b) rejects Elimination insofar as rationality does not consist in responding to reasons one has. However, Kolodny thinks that reasonable agents are necessarily consistent. See Buchak and Pettit (2015), Guindon (2014; 2016) and Reisner (2011) for discussion. See also Kolodny (2008a; 2008b) and Raz (2005, 6) for discussion related to the practical realm.

Elimination. Necessarily, if an agent takes inconsistent attitudes towards P at time t, then he or she violates a substantive requirement of epistemic rationality other than Consistency. In other words, Consistency plays no distinct explanatory role in theories of epistemic rationality.

A quick clarificatory remark: I here insist on the notion of "*distinct* explanatory role." Violating Consistency could be part of an explanation of why, in a given situation, an agent is epistemically irrational without being the *only* explanation of why an agent is epistemically irrational. For example, in some situations where an agent is inconsistent, there could be multiple sufficient explanations of why he or she is irrational.[114] Still, I am not interested in putative cases where an agent's irrationality is overdetermined. As long as inconsistent agents violate requirements of rationality other than Consistency, we can provide an explanation of why they are irrational without referring to Consistency, which supports Elimination.

Elimination would be sufficient to vindicate the Modest Reductionist Hypothesis. Recall that, in chapter 3, I argued for the following: when compared with substantive requirements of epistemic rationality, putative requirements such as Inter-Level Coherence or Intra-Level Coherence have no distinct explanatory role in the theory of epistemic rationality. This left us with one possibility to analyze, namely, that Consistency plays a distinct explanatory role. In such a context, if Elimination were true, the Modest Reductionist Hypothesis would be vindicated.

The Strong Reasons-Responsiveness Thesis entails Elimination. Here is why. Defenders of the Strong Reasons-Responsiveness Thesis think that there are other requirements of epistemic rationality susceptible of vindicating Elimination, such as responding correctly to epistemic reasons one has. According to the Strong Reasons-Responsiveness Thesis, epistemic reasons can't simultaneously warrant the belief that P and the belief that ~P. This means that one is never simultaneously permitted to believe P and to disbelieve P. So, if one simultaneously believes P and ~P, one has necessarily failed to respond to one's epistemic reasons. Hence, Elimination is entailed by the Strong Reasons-Responsiveness Thesis.

[114]Fogal (m.s.) endorses such a view.

The explanatory role played by Consistency is connected to broader debates concerning the normativity of epistemic rationality. Recall the objection from truth-conduciveness discussed in section 1.2.5: according to Kolodny, what matters from an epistemic point of view is acquiring true beliefs and avoiding false beliefs. However, satisfying Consistency does not guarantee a better ratio of true to false beliefs. So, except if Consistency doesn't play a distinct explanatory role in the theory of epistemic rationality, we will have to answer Kolodny's challenge.[115]

4.1.2. Can the Normativity of Consistency Be Vindicated?

In chapter 1, I have presented many objections against the normativity of structural requirements of epistemic rationality. Some of these objections have to do with bootstrapping, others have to do with truth-conduciveness, and so forth. In view of these objections, those who think that epistemic rationality is normative have two options: they can provide a new explanation of the normativity of structural requirements and answer the aforementioned challenges, or they can argue that Consistency plays no distinct explanatory role in the theory of epistemic rationality. This chapter will not try to argue that Consistency is a normative requirement of epistemic rationality. My strategy is different: I argue that when it comes to ideally rational agents, Consistency doesn't play a distinct explanatory role in the theory of epistemic rationality. In other words, I take the second option.

Yet, it should be noted that a case can be made in favour of the normativity of Consistency. Recently, a new strategy has emerged to vindicate the normativity of Consistency. It comes from decision theory and relies on the following principle:

Strong Accuracy-Dominance. If an available state X is strongly accuracy-dominated by an available state Y at every possible world, in the sense that state Y is epistemically better or is more accurate than state X at every possible world, one ought to avoid state X.

[115]Some deny that epistemic norms have to do with the goals of getting true beliefs and avoiding false beliefs. See the debate between Berker (2013a; 2013b; 2015) and Goldman (1986; 2015). Others admit that explaining the normativity of Consistency is an important challenge. See, for instance, Broome (2008; 2013, chap. 11) and Way (2009; 2010).

There is a weak and a strong interpretation of what is entailed by the accuracy-dominance arguments. According to the strong interpretation, accuracy-dominance arguments entail that one ought to be consistent. Joyce, for instance, argues that:

> It is thus established that degrees of belief that violate the laws of probability are invariably less accurate than they could be. Given that an epistemically rational agent will always strive to hold partial beliefs that are as accurate as possible, this vindicates the fundamental dogma of probabilism [according to which degrees of belief must make conformity to the axioms of probability]. (Joyce 1998, 600)

According to the weak interpretation, accuracy-dominance arguments merely entail that one ought not to be inconsistent. Easwaran, for instance, indicates that "we can use dominance to eliminate" the inconsistent doxastic options (Easwaran 2015, 826). In other words, dominance is here used to argue against inconsistency. Thus, we can make the following distinction between the two views:

Normativity+ of Consistency. Given the accuracy-dominance arguments, A ought to be consistent.

Normativity- of Consistency. Given the accuracy-dominance arguments, A ought not to be inconsistent.[116]

Plausibly, Strong Accuracy-Dominance can vindicate Normativity-. Here is why. An agent can take different doxastic attitudes towards P, as in the following:

(i) Believing P and not disbelieving P;

(ii) Disbelieving P and not believing P;

(iii) Neither believing nor disbelieving P;

(iv) Believing P and disbelieving P.

[116]It may seem weird to make a distinction between reasons to be consistent and reasons not to be inconsistent. However, some philosophers such as Snedegar (2018) have argued that we should make a distinction between reasons for P and reasons against ~P. He roughly argues that the problem with views that assimilate reasons against to reasons for is that there can be reason not to go for P, even if there are worse alternatives to P. For instance, suppose that I am trying to decide what to drink. I might have conclusive reason not to drink gin, but this doesn't entail that I have a reason to drink any beverage that isn't gin. I should definitely not drink petrol, even if petrol isn't gin. This is entirely compatible with my having conclusive reason not to drink gin. The same could be said of reasons to be consistent and reasons not to be inconsistent. I discuss this argument in Daoust (m.s. c).

The question is whether (iv) is strongly dominated. To answer this question, we need to determine the epistemic value of (iv) at every possible world. In a veritist framework, only true beliefs have final epistemic value and only false beliefs have final epistemic disvalue. Accordingly, T is the epistemic value of having a true belief (for T>0), F is the epistemic disvalue of having a false belief (for F<0), and the epistemic value of not believing P (or not disbelieving P) is 0.[117] Finally, assume that T<-F, which amounts to endorsing a conservative account of epistemic value (I will come back to this point in section 4.3.1). Such a conservative constraint is extremely plausible. Indeed, as Dorst explains:[118]

> [An epistemically rational agent] will be doxastically conservative... Why? Well here's a fair coin—does she believe it'll land heads? Or tails? Or both? Or neither? Clearly neither. But if she cared more about seeking truth than avoiding error, why not believe both? She'd then be guaranteed to get one truth and one falsehood, and so be more accurate than if she believed neither... Upshot: we impose a *Conservativeness* constraint to capture the sense in which Rachael has 'more to lose' in forming a belief than she does to gain. (Dorst 2017, 11)

In view of the foregoing, we can determine the possible value of each option at every possible world. Since the value of options is solely determined by P's truth value, we need to consider the worlds in which P is true and the worlds in which P is false, as in the following table:

Table 4.1. Doxastic options towards P agents have

Doxastic options / possible world	P is true	P is false
Believing P and not disbelieving P	T	F
Disbelieving P and not believing P	F	T
Neither believing nor disbelieving P	0	0
Believing P and disbelieving P	T+F	T+F

[117] I'm glossing over some inessential subtleties here. It is possible to assign a value to not believing P (or to withholding judgment on whether P), but ultimately, we would get exactly the same results. See Easwaran (2015, §C) and Dorst (2017, 10 n12).

[118] In addition to Dorst's argument, see Easwaran (2015), Easwaran and Fitelson (2015) and Pettigrew (2016b) for similar arguments in favour of the conservative account of epistemic value. See also Steinberger (forthcoming), who argues that the conservative constraint is not derived from considerations of accuracy alone.

Finally, in accordance with table 4.1, we can conclude that inconsistent combinations of beliefs are strongly dominated. The following reasoning supports such a conclusion:

(1) T<-F (assumption). Accordingly, T+F<0.

(2) Following (1) and table 4.1, believing P and disbelieving P simultaneously has an epistemic value of less than 0 at every possible world.

(3) However, following table 4.1, neither believing nor disbelieving P has an epistemic value of 0 at every possible world.

(C) Therefore, following (2) and (3), inconsistent combinations of beliefs such as believing P and disbelieving P are strongly dominated: another available option (neither believing nor disbelieving P) is more valuable at every possible world.[119]

Hence, if Strong Dominance and the conservative account of epistemic value are correct, it is plausible that one ought to avoid being inconsistent. In other words, Strong Dominance and Conservativeness entail Normativity-. Minimally, this suggests that there can be weak vindications of the normativity of Consistency.[120] But as I explained at the beginning of this section, I here favour a different strategy, which is to show that Consistency doesn't play an explanatory role in theories of epistemic rationality for ideal agents. So, I leave aside the accuracy-dominance strategy.

4.2. A Blind Spot in the Strong Reasons-Responsiveness Thesis

I will now argue against the Strong Reasons-Responsiveness Thesis. Defenders of such a view assume that, since epistemic rationality consists in responding correctly to epistemic reasons one has, Consistency plays no distinct explanatory role in theories of epistemic rationality. I will argue that such an assumption is mistaken. The fact that epistemic rationality has to do with responding correctly to epistemic reasons one has merely entails that Consistency plays no distinct explanatory role *in cases where epistemic reasons are impermissive*. So, defenders of the Strong Reasons-Responsiveness Thesis assume that a specific type of permissiveness is false, which is far from obvious.

[119]Similar arguments can be found in Easwaran (2015, §B) and Pettigrew (2016b, 256). Dorst (2017, 31—esp. proposition 3) argues for a similar but contextualist view.
[120]I discuss this view in Daoust (m.s. c).

Here is how I will proceed. First, I will introduce Reasons Permissiveness and explain why there is an essential connection between the extreme version of Permissiveness and Elimination. Then, I will explain why the possibility of permissive epistemic reasons is a blind spot in the Strong Reasons-Responsiveness Thesis.

4.2.1. Moderate and Extreme Reasons Permissiveness

Evidential Permissiveness roughly states that, relative to a given body of evidence, it can be epistemically rational for A to believe P and not to believe P.[121] While this chapter is strongly inspired by the literature on Evidential Permissiveness, I am not interested in the normative status of Evidentialism. So instead of focusing on an agent's evidence, I will rather focus on an agent's epistemic reasons.[122] Reasons permissiveness is to epistemic reasons what evidential permissiveness is to evidence. In view of the foregoing, Reasons Permissiveness would roughly state that, relative to a body of epistemic reasons, it can be epistemically rational for A to believe P and not to believe P.

As noted by White, there is a moderate and an extreme sense in which evidence can be permissive (White 2005, secs. 2–4). Relative to a body of evidence, it could be rational for one to believe P and also *not to believe* P, but it could also be rational for one to believe P and to *disbelieve* P, which is logically stronger. Indeed, it is possible not to believe P while not disbelieving P (notably by withholding judgment on whether P). Following White's distinction between moderate and extreme Evidential Permissiveness, we can distinguish moderate and extreme versions of Reasons Permissiveness, as follows:[123]

Moderate Reasons Permissiveness. Relative to a body of epistemic reasons, A can be rationally permitted to believe P and not to believe P.

[121]See Kopec and Titelbaum (2016), White (2005; 2014) and Kelly (2014) and for an overview of the debate surrounding the evidential interpretations of Permissiveness.
[122]See Littlejohn (2012) or Owens (2002) on why there could be a distinction between epistemic reasons and evidence. I here remain neutral on whether such a distinction is correct. And again, in this chapter, I refer to epistemic reasons in the broad sense, which includes facts about appearances.
[123]Obviously, since evidence appears to be the main type of epistemic reason, there is a close connection between Reasons Permissiveness and Evidential Permissiveness. Specifically, if epistemic reasons are permissive, this probably means that evidence is permissive. Nevertheless, I prefer to distinguish the two views and focus on the former.

Extreme Reasons Permissiveness. Relative to a body of epistemic reasons, A can be rationally permitted to believe P and to believe ~P.

I will soon argue that *Extreme* Reasons Permissiveness holds if and only if Elimination is false. For now, I will explain why Moderate Reasons Permissiveness is compatible with Elimination.

Moderate versions of Permissiveness are inconclusive for establishing the explanatory role of Consistency. Consistency is about maintaining specific combination of attitudes. Moderate reasons permissiveness merely permits an agent to believe P and not to hold this attitude. However, an attitude and the *absence* of an attitude cannot be jointly inconsistent, since consistency governs combinations of attitudes. Since I am interested in situations where an agent is rationally permitted to hold jointly inconsistent attitudes, moderately permissive situations are not sufficient for rejecting Elimination.

Furthermore, Consistency could not play any explanatory role here, since it is *metaphysically* (or logically) impossible for an agent to believe P and not to believe P simultaneously. Given this, moderately permissive situations cannot help to establish that Elimination is false. Here is why. Recall that Elimination is a conditional: if an agent holds incompatible attitudes towards P at time t, then he or she violates a requirement of epistemic rationality other than Consistency. However, since it is *metaphysically impossible* for an agent to believe P and not to believe P simultaneously, this means that the antecedent of Elimination will never obtain in moderate permissive situations. In other words, since it is impossible to believe P and not to believe P simultaneously, it is impossible to satisfy the antecedent of Elimination with such a combination of attitudes. If the antecedent of a conditional is always false, then the conditional is trivially true. So again, moderately permissive situations are not sufficient for rejecting Elimination.

4.2.2. Reasons Permissiveness and Elimination

I will now argue that Elimination is true if and only if Extreme Reasons Permissiveness is false.

Why is there such a connection between Extreme Reasons Permissiveness and Elimination? To see how these two theses are related, let's analyze a potential argument against Extreme Reasons Permissiveness. Suppose that there are situations in which one is permitted to believe P and also to disbelieve P. If, relative to a body of epistemic reasons, an epistemically rational agent is permitted to believe P and also permitted to believe ~P, then nothing could prevent him or her from believing P and believing ~P. In other words, if two incompatible attitudes are rationally permitted, it seems that nothing forbids an agent from believing both propositions simultaneously. Hence, if there is a constraint prohibiting agents from believing incompatible attitudes, it must come from the fact that Extreme Permissiveness is false.

Of course, the above line of reasoning is inconclusive. Rather than following from Permissiveness, the above conclusion comes from an invalid modal inference. As Ballantyne and Coffman (2011, 9) point out, the modal scope of permissions must be taken into account. Specifically, assuming that ◊ is the modal operator for permissions, ◊[B(P)^B(~P)] is not entailed by [◊B(P)^◊B(~P)]. For example, if you are permitted to drink and you are permitted to drive, it doesn't follow from such permissions that you are permitted to drink *and* drive. While an agent may have different permissions, this doesn't mean that all of his or her permissions can be satisfied in the same possible world. Therefore, while an agent could be permitted to believe P and to believe ~P, it doesn't follow that he or she is permitted to have both beliefs simultaneously.

If ◊B(P) and ◊B(~P) are only rationally accessible in mutually exclusive worlds, this means that there is a rational requirement in every possible world prohibiting the conjunction [B(P)^B(~P)]. Such a requirement would be Consistency. However, assuming that epistemic reasons can be permissive in some possible worlds, such a requirement cannot be a consequence of an agent's epistemic reasons. In other words, in permissive situations, Consistency governs an agent's attitudes even if the epistemic reasons one has may warrant believing P and believing ~P. Thus, if Extreme Permissiveness is true, a requirement like Consistency cannot be reduced to reasons-responsiveness requirements, since incompatible doxastic attitudes are not ruled out by the epistemic reasons one has.

In view of the foregoing, Elimination entails that epistemically permissive situations (of the extreme sort) are impossible. In a permissive picture of epistemic rationality, Consistency plays a distinct explanatory role. First, assume that Extreme Reasons Permissiveness is true. If, relative to a body of epistemic reasons, distinct incompatible attitudes towards P are rationally permitted, we need a requirement like Consistency to prohibit an agent from believing P and believing ~P simultaneously. Even if epistemic reasons warrant distinct incompatible beliefs, Consistency will prohibit an agent from believing contradictory propositions. In other words, if epistemic reasons are permissive, then Consistency has a genuine explanatory role and Elimination is false. Now, assume that Extreme Reasons Permissiveness is false. If epistemic reasons are impermissive (or moderately permissive), then an agent's epistemic reasons never warrant believing P and believing ~P. This means that an agent should *never* end up believing P and ~P on the basis of his or her epistemic reasons. So, Consistency would have no distinct explanatory power and Elimination would be true.

4.2.3. The Possibility of Extreme Reasons Permissiveness Undermines the Strong Reasons-Responsiveness Thesis

Let's assume for a moment that Extreme Reasons Permissiveness is a live possibility (in section 4.3, I will explain why such a view should be taken seriously). I previously argued that epistemic rationality has to do with responding correctly to epistemic reasons one has, or, at least, that making a separation between reasons and rationality leads to undesirable consequences. However, this doesn't mean that Consistency plays no distinct explanatory role in theories of epistemic rationality. In cases where epistemic reasons warrant believing P and believing ~P, Consistency plays a distinct explanatory role, since an epistemically rational agent should not believe P and believe ~P simultaneously. So, the fact that rationality has to do with responding to reasons one has does not entail that Elimination is true, since Extreme Reasons Permissiveness could be true. Defenders of the Strong Reasons-Responsiveness Thesis implicitly assume that epistemic reasons are impermissive (or, at least, no more than moderately permissive).

This leads me to conclude that we should endorse the Modest Reasons-Responsiveness Thesis. According to such a view, Consistency plays no distinct explanatory role *in impermissive situations*. However, in cases where both believing P and disbelieving P are warranted by a body of epistemic reasons, Consistency plays a distinct explanatory role.

An implication of my argument is that the problem surrounding the normativity of Consistency can affect any theory of epistemic rationality. Whether we think that rationality has to do with structural requirements or with substantive requirements, Consistency can play a distinct explanatory role in theories of epistemic rationality. The only way to deny this possibility is to argue against Extreme Reasons Permissiveness. Otherwise, Kolodny's teleological argument against the normativity of Consistency (discussed in section 1.2.5) will also apply to substantive theories of epistemic rationality. Under the assumption that Extreme Reasons Permissiveness is true, responding to epistemic reasons one has "may as soon lead one away from, as toward, the true and the good" (Kolodny 2007b, 231).

4.3. Permissiveness and Kolodny's Decision-Theoretic Argument

Many philosophers will think that the argument of this chapter is uninformative, notably because they are convinced that epistemic reasons never warrant the belief that P and the belief that ~P simultaneously. For instance, it is often argued that one should not believe ~P if one has sufficient epistemic reason to believe P, and that one should not believe P if one lacks sufficient epistemic reason to believe P. Since epistemic reasons for believing P are either sufficient or insufficient, it appears trivially true that one is never simultaneously permitted to believe P and to disbelieve P. Hence, it seems that Extreme Reasons Permissiveness is false (or that Elimination is correct).[124]

In this section, I will argue that such claims about the sufficiency of epistemic reasons are ambiguous or misleading: they presuppose that the weight of epistemic reasons is objectively determined. In order to show this, I will offer an extended reconstruction of

[124]Some authors take such a line of reasoning to be obvious. See notably Lord (2017, 9), Kiesewetter (2017, 180-85), Matheson (2011) or Sylvan (2015b). See also Foley's (1987, chap. 6.2) argument against the epistemic rationality of believing contradictory propositions.

Kolodny's (2007b, 232-37) decision-theoretic argument for the claim that a body of epistemic reasons never supports believing P and disbelieving P simultaneously.[125] Then, I will explain why such an argument misses the target and the significance of such a result in the debate surrounding the explanatory role of Consistency.

4.3.1. Kolodny's Argument from Epistemic Decision Theory

According to a teleological perspective like Kolodny's, one ought to optimize the balance of wanted events and unwanted events. When making rational decisions under uncertainty, one ought to consider how various decisions might turn out and what the risk associated with each of these potential decisions is. In the epistemic realm, risk can be understood as the probability of forming an epistemically unwanted belief—that is, a false belief. Hence, to rationally believe P under uncertainty means that believing P is an *optimal* epistemic decision with respect to epistemic risk.

Following Kolodny, let's assume that epistemically rational agents ought to maximize *expected* epistemic utility.[126] In such a context, determining what one is rationally permitted to believe amounts to an optimization problem. Suppose that Z is the epistemic probability or the rational credence in P, for $0 \leq Z \leq 1$.[127] Suppose that F is the epistemic value of falsely believing that P, and that T is the epistemic value of truly believing that P.[128] Let's also assume the following constraints on the values of F and T:

---

[125]There are two reasons why I here offer an extended version of the argument. First, Kolodny's original argument is stated very quickly. Second, since the publication of Kolodny's argument, Easwaran (2015), Pettigrew (2016c) and Dorst (2017) have developed similar frameworks that are much more comprehensive. In such a context, I prefer to develop an extended reconstruction of Kolodny's argument. This allows me to consider the strongest interpretation of his view.

[126]The expected value is sometimes called the weighted mean value. For example, suppose that, in a fair lottery, 10 participants each have 1 chance in 10 of winning a single prize of $20. In that lottery, 9 participants won't win anything, and 1 participant will win $20. Since $(9 \cdot 0 + 1 \cdot 20)/10 = 2$, the weighted mean value of this lottery is $2. This means $2 is the expected prize for each participant. See Buchak (2013) for alternatives to expected utility theory.

[127]Kolodny's argument has to do with epistemic probabilities (Kolodny 2007b, 233). However, under the assumption that rational credences track epistemic probabilities, these two notions can be used nearly interchangeably.

[128]Kolodny considers the epistemic value of truly believing P and the epistemic value of avoiding a false belief concerning P. Following Easwaran (2015), Pettigrew (2016c) and Dorst (2017), I take my framework to be more intuitive. In any case, nothing hinges on this small difference.

(i) $T > 0$

(ii) $F < 0$

Finally, suppose that the epistemic value of not forming a belief about P is 0.[129] When an epistemically rational agent forms a belief about P, the following two rules must be satisfied:[130]

(iii) If $(0 < Z \cdot T + (1-Z) \cdot F)$, then taking the epistemic risk of believing that P is permitted, since the expected epistemic value of believing P is uniquely optimal.

(iv) If $(0 > Z \cdot T + (1-Z) \cdot F)$, then taking the epistemic risk of believing that P is not permitted, since the expected epistemic value of believing P is suboptimal.

An example might be helpful here. Suppose that an agent's rational credence in P is 0.4, that the epistemic value of forming a false belief about P is -5, and that the epistemic value of forming a true belief about P is 2. In such a case, $Z \cdot T + (1-Z) \cdot F = (0.4 \cdot 2) + (0.6 \cdot -5) = -2.2$. This means that the expected epistemic utility associated with believing P is -2.2. Since $0 > -2.2$, the epistemic value of not believing that P (0) is higher than the expected epistemic value of believing that P (-2.2). So far, this means that the agent should not believe P. Now, let's calculate the expected epistemic value of believing that ~P. If an agent's rational credence in P is 0.4, then his or her rational credence in ~P is 0.6. Again, let's assume that the epistemic value of forming a false belief about P is -5, and that the epistemic value of forming a true belief about P is 2. This means that $Z \cdot T + (1-Z) \cdot F = (0.6 \cdot 2) + (0.4 \cdot -5) = -0.8$. Since $0 > -0.8$, the expected epistemic value of not believing that ~P (0) is higher than the expected epistemic value of believing that ~P (-0.8). Therefore, in such a situation, the agent should not believe that P and not believe that ~P, which amounts to suspending judgment on whether P. That is, the uniquely optimal option is to withhold judgment on whether P.

---

[129] Zero is simply a reference point. Suppose that, in order to decide whether to believe P, an agent calculates the expected utility to believe that P. Suppose that the result is -10. Since the epistemic value of not believing that P is 0 *by reference*, this means that there are 10 utiles associated with not believing that P. When the expected value of forming a belief that P is under 0, this means that not believing that P is a better epistemic option *with reference to an epistemic value of 0*. Dorst (2017, 9-12) makes similar remarks.

[130] It should be noted that similar principles have been developed elsewhere since the publication of Kolodny's argument. See, for instance, Easwaran (2015), Pettigrew (2016c) and Dorst (2017). Following Kolodny, I here assume that rational beliefs are determined by epistemic probabilities or rational credences. By contrast, Easwaran uses these principles to argue that rational credences are determined by rational beliefs.

Now, are there cases where the expected epistemic value of believing P is equal to the expected epistemic value of believing ~P? If there were such cases, they would entail that Extreme Reasons Permissiveness is correct. Consider, for example, a case where the rational credence in P is 0.5 and the rational credence in ~P is also 0.5. According to Kolodny, such cases vindicate Extreme Reasons Permissiveness as long as we assume that -F=T. Under such an assumption, we get the following result:

(i) If one is rationally permitted to believe P and to disbelieve P, this means that $0 = T \cdot Z + F \cdot (1-Z) = T \cdot (1-Z) + F \cdot Z$
(ii) Assume that the rational credence in P is 0.5 and that the rational credence in ~P is 0.5. Following (i), $0 = T \cdot 0.5 + F \cdot 0.5 = T \cdot 0.5 + F \cdot 0.5$.
(iii) Following (ii), 0 = T+F. So, -F=T. Hence, -F=T if it can be equally optimal to believe P and to disbelieve P.

However, under the assumption that -F=T, one is *never* rationally required to suspend judgment on whether P.[131] That is, for any rational credence in P, either one is rationally permitted to believe P or one is rationally permitted to disbelieve P. But such a conclusion is untenable. There are cases where withholding judgment on whether P is the only rational option. So, it seems perfectly plausible to assume that T<-F. This amounts to endorsing what Kolodny calls the "conservative" account of epistemic value (Kolodny 2007b, 234).

According to Kolodny, the constraint T<-F entails that Extreme Reasons Permissiveness is false. Indeed, relative to variables Z, F and T, it is never equally optimal to believe P and to disbelieve P. Hence, Kolodny concludes that "in any given situation, either it will be the case that (one is required by reason not to believe p), or it will be the case that (one is required by reason not to believe not-p)" (Kolodny 2007b, 236). In other words, Elimination seems to be vindicated.

4.3.2. Why Kolodny's Argument Is Problematic

In accordance with Kolodny, I believe that the kind of decision-theoretic model discussed in the previous subsection is interesting. I also believe that T<-F, the constraint on the value of having true beliefs and the value of having false beliefs, is very plausible. We can

[131] As I indicated in section 4.1.2, Dorst (2017, 11) and Easwaran (2015, 824) reach similar conclusions.

even assume that the values of T and F are uniquely determined on an agent's body of epistemic reasons.[132] The fundamental problem with Kolodny's argument concerns the assumption that epistemic probabilities or rational credences, which represent the weight of one's epistemic reasons, are uniquely determined. It is far from obvious that epistemic probabilities are uniquely determined. Furthermore, without such an assumption, Kolodny's argument collapses.

According to many philosophers, epistemic reasons are sufficient or insufficient to the extend that they have been mediated through a set of *epistemic standards*. An agent's epistemic standards are the rules, models or assumptions he or she relies on to evaluate epistemic reasons. Such a notion can be understood in a broad sense, including background beliefs, standards of reasoning, prior probability distributions and the like.

Now, there can be distinct incompatible epistemic standards one can entertain. Accordingly, this means that an agent can have sufficient epistemic reason to believe P *relative to standard A*, but he or she can also have sufficient epistemic reason to disbelieve P *relative to standard B*. If this is correct, Kolodny's decision-theoretic argument fails.

Permissiveness concerning epistemic standards can take many forms. First, epistemic probabilities could be permissive, or there could be more than one rational credence in P an agent can entertain. Such a possibility would affect whether an agent has sufficient epistemic reason to believe P. Meacham (2014), for example, has argued that there is more than one rational credence function an epistemically rational agent can entertain. For example, perhaps Linda is rational in believing P if and only if she is rational in having a credence of more than 0.6 in P. However, suppose that epistemic probabilities are permissive, and that it is rational for Linda to entertain any credence in P located in the interval [0.3, 0.7]. If she entertains a credence of 0.65 in P, it will be rational for her to believe P, but if she entertains a credence of 0.65 in ~P (or a credence of 0.35 in P), it will be rational for her to disbelieve P. Thus,

[132]It should be noted that such a claim is not uncontroversial. For example, jamesian pragmatists think that there is no uniquely rational epistemic value of true beliefs and epistemic disvalue of false beliefs. See Kelly (2014, sec. 2), Pettigrew (2016) and Dorst (2017) for discussion.

assuming that sufficiency has to do with rational credences, such a notion is not necessarily uniquely determined.

According to Meacham, credal permissiveness is entailed by the fact that we can't identify a single credal function for ignorant agents. Impermissive Bayesians think that all rational ignorant agents start their credal lives with the same prior credence function (that is, the credence function one has before acquiring evidence). But as noted by Meacham, it is hard to see why there would be such demanding constraints on an agent's priors. Various principles governing rational priors have been proposed, but they are not impermissive (Meacham 2014, 1192-3). Take the Principal Principle, which roughly states that an agent's knowledge of the objective probability that P constrain his or her credence in P. The Principal Principle tells us which credence in P we ought to entertain when we know P's objective probability, but it doesn't tell us which credence in P to entertain when we ignore P's objective probability. So, the Principal Principle leaves some slack.

There is one putative impermissive principle governing prior credences, namely, the Principle of Indifference. According to this principle, if an agent has no evidence for or against the various propositions in a partition containing n elements, then he or she should entertain a credence of 1/n in each proposition. However, according to Meacham, defenders of the Principle of Indifference face a trilemma: the principle is either trivial, inconsistent or arbitrary (Meacham 2014, 1201-2).

Of course, one could revise Kolodny's framework in order to avoid representing the weight of epistemic reasons with epistemic probabilities. However, any representation of the weight of epistemic reasons will eventually face the same difficulties. Titelbaum and Kopec (forthcoming; m.s.), Goldman (2010), Meacham (2014), Schoenfield (2014) and Sharadin (2015) have argued that epistemically rational agents can entertain incompatible epistemic standards, regardless of how epistemic reasons are represented. According to them, there is no objective measure of the weight of epistemic reasons. Let's call such a view the Permissive Epistemic Standards Thesis, as in the following:

Permissive Epistemic Standards Thesis. The weight of epistemic reasons in favour of P is not objectively determined—rather, in order to determine their weight, epistemic reasons have to be subjectively mediated through an agent's rational epistemic standards. Furthermore, there are multiple incompatible rational epistemic standards, in the sense that there are incompatible but equally rational ways to draw various conclusions from a body of epistemic reasons.

For example, in the process of evaluating whether he or she should believe a proposition, an agent may assign different weights to various epistemic reasons. So, an agent may reach incompatible conclusions depending on how he or she is weighting the epistemic reasons. Hence, there could be no unique notion of sufficient epistemic reason or unique way to reason from a body of epistemic reasons. Presupposing that the notion of sufficiency is objectively determined obliterates a plausible type of permissiveness.

There are two main arguments that have been offered in favour of the Permissive Epistemic Standards Thesis. According to Schoenfield, epistemic standards are not like any other belief. They are the considerations in virtue of which we evaluate epistemic reasons. This entails that we cannot provide independent justification in favour of our rational epistemic standards, since those standards are precisely the considerations in virtue of which we evaluate our doxastic states.[133] She argues:

> We can never give reasons for why we weigh the evidence in one way rather than another that are independent of everything else. This is just a fact about epistemic life that we have to live with: the methods that we use to evaluate evidence are not the sorts of things we can give independent justification for. (Schoenfield 2014, 202)

Accordingly, she thinks that asking for independent justification in favour of our epistemic standards would lead to widespread skepticism. Indeed, agents could not provide such a justification, and thus would have to stop entertaining any standard. Schoenfield takes such a conclusion to be absurd. Then, the only option left is to allow for rational epistemic standards that are not supported by independent justification. This supports the Permissive Epistemic Standards Thesis.

[133]This line of reasoning is similar in fashion to the argument against the possibility of rationally evaluating hinge propositions. See Coliva (2015), Pritchard (2016) and Siegel (2019) on this debate.

Goldman, Titelbaum and Kopec offer a different argument in favour of the Permissive Epistemic Standards Thesis. According to them, epistemic standards are rational insofar as they meet some criteria, such as being internally consistent and sufficiently reliable. However, distinct incompatible standards can be consistent and sufficiently reliable. So, equally rational agents could entertain distinct incompatible epistemic standards, which supports the Permissive Epistemic Standards Thesis (I take a closer look at this argument in the next chapter).

In summary, Kolodny thinks that Elimination is correct because he assumes that the weight of epistemic reasons is objectively determined. But such an assumption is far from trivial or unproblematic. More importantly, such an argument is unsuccessful against permissiveness, since it presupposes that the weight of epistemic reasons is uniquely determined.

4.3.3. Refining the Modest Reasons-Responsiveness Thesis

I will now explain why the Permissive Epistemic Standards Thesis sheds a new light on Elimination.

The Permissive Epistemic Standards Thesis reveals that Consistency could play an unexpected explanatory role. Surely, relative to a body of evidence, a credence function, and a set of epistemic standards, it is highly plausible that there is a unique rational attitude to hold towards P. Still, this does not vindicate Elimination. Consistency could have an explanatory power by prohibiting some combinations of epistemic standards an agent can entertain. In other words, as long as distinct incompatible epistemic standards are rationally permitted, Consistency plays an explanatory role by prohibiting some combinations epistemic standards.

In view of the foregoing, here is how we can refine the Modest Reasons-Responsiveness Thesis. Rationality consists in part in responding to reasons one has. However, epistemic reasons support a doxastic attitude towards P only insofar as they have been subjectively mediated through an agent's epistemic standards. If an agent is rationally permitted to entertain distinct incompatible epistemic standards, Consistency plays a distinct

explanatory role in governing the combinations of epistemic standards one can entertain. But in cases where epistemic standards are uniquely determined,[134] or in cases where epistemic reasons are impermissive, Consistency plays no distinct explanatory role.

Here is an example of how Consistency could play an explanatory role in the combinations of epistemic standards one can entertain. Suppose that an agent can make a choice between two sets of epistemic standards: A and B. For the sake of the argument, let's assume that standards A and B are equally good. For instance, we can presume that these sets of standards are equally reliable: relative to a body of evidence, satisfying A leads one to the right answer 90% of the time and satisfying B also leads one to the right answer 90% of the time. We can also assume that the agent does not have the impression that one set is more plausible or commonsensical than the other. In such a situation, it seems plausible that it is rational for an agent to entertain standards A, but also to entertain standards B. However, since A and B are incompatible with each other, an agent cannot entertain both sets of standards A and B simultaneously (it would be inconsistent for him or her to entertain both sets of standards). In such a case, Consistency plays a distinct explanatory role: the reason why an epistemically rational agent should not entertain both standards A and B simultaneously is that such a combination of standards violates Consistency.

4.4. Extending The Explanatory Role of Consistency to Collective Contexts

4.4.1. The Distinction Between Intrapersonal and Interpersonal Permissiveness

So far, I have focused on the case of one agent (e.g., on intrapersonal cases). For instance, I have focused on the possibility that one agent is epistemically permitted to entertain distinct incompatible standards. We can weaken this assumption and take into account

[134]One could wonder how some epistemic standards could be uniquely determined. For instance, some epistemic standards could be subject to an "overlapping consensus"—that is, the various rational epistemic systems could concur that some specific epistemic standards are correct, even if the grounds in favour of such standards may differ from one system to another. Alternatively, it could be argued that there are self-justified standards. For example, it is plausible that a standard roughly stating "trust your direct perceptions" is self-justified.

collective contexts in which distinct agents entertain distinct epistemic standards. For instance, consider the following case discussed by Miriam Schoenfield (2014):

> I was once talking to a very religious friend... about whether or not her particular religious beliefs were justified. As these conversations tend to go, we each proposed arguments that challenged the other's beliefs, responded to them, deemed the other's responses unsatisfactory, and neither of us budged (Schoenfield 2014, 193).

Schoenfield describes a case of *persistent* disagreement between two epistemic peers: even when disclosing their arguments, evidence and objections, they stand their ground. And assuming that these agents are unable to settle their dispute by presenting their evidence and reasoning (which they disclose to each other), we have a case of *deep* disagreement between epistemic peers.

There is a simple explanation of why two agents (call them Kate and Brad) who disclose their evidence and reasoning sometimes persistently disagree about some claim: one of them could be irrational. However, Schoenfield offers a different explanation. She thinks that epistemically rational agents can find themselves in a *permissive situation*, in the sense that, relative to the same body of evidence, it is rational for one agent to believe P and it is rational for the other to disbelieve P. However, Schoenfield goes a step further and argues that Kate can be rationally *required* to believe P while Brad is rationally *required* to disbelieve P. So, even when disclosing their evidence, reasoning and objections, they *ought* to disagree. Hence, if Schoenfield is right, there are cases of fundamental and persistent rational disagreement between epistemic peers.

In order to make sense of Schoenfield's view, we need to make a distinction between intrapersonal and interpersonal versions of permissiveness, as in the following:

Intrapersonal Permissiveness. Relative to a body of epistemic reasons, A can be rationally permitted to believe P and to believe ~P.

Interpersonal Permissiveness. Relative to a body of epistemic reasons, two agents can be rationally permitted to entertain distinct incompatible beliefs. For example, agent 1 could be rationally permitted to believe P and agent 2 could be rationally permitted to believe ~P.

Relative to the above distinction, there are three possible stands concerning epistemic permissiveness. Some authors have argued that permissiveness is false *both* at the intrapersonal and at the interpersonal level,[135] or that permissiveness is true *both* at the intrapersonal and at the interpersonal level.[136] Others (like Schoenfield) endorse both intrapersonal uniqueness and interpersonal permissiveness.[137] If such a compromise solution is correct, there are situations in which, relative to a body of evidence shared by two agents, there is no uniquely rational answer at the interpersonal level, but one agent is required to believe P and the other is required to disbelieve P.

4.4.2. Interpersonal Permissiveness in Collective Contexts

If interpersonal permissiveness is true, Consistency can play an explanatory role in collective contexts. Buchak and Pettit's (2015) argument from Group Coherence, which analyzes the explanatory role of Consistency in collective contexts, provides a good illustration of this.

How can we determine which beliefs are held by group agents? Presumably, a rule like aggregation can indicate the group agent's doxastic states. Thus if the majority of agents that are part of a collective believes P, it seems that the group agent believes P. Now, how does group agents in epistemically permissive situations show that Elimination is false? According to Buchak and Pettit, it is possible that agents who share the same epistemic reasons and who responded correctly to them form beliefs that are jointly inconsistent (Buchak and Pettit 2015, 215), as in the following:

---

[135]See Dogramaci and Horowitz (2016), D. Greco and Hedden (2016), Hedden (2015b), Horowitz (2014b) and Matheson (2011). White (2005; 2014) also argued for interpersonal uniqueness, but many authors suggested that his arguments support only intrapersonal uniqueness (see Kopec and Titelbaum (2016) on this specific issue).

[136]See Drake (2016), Douven (2009), Titelbaum and Kopec (forthcoming; m.s.), Kopec (2015) and Raleigh (2015). Brueckner and Bundy (2012) reject White's argument for uniqueness, but they do not endorse a specific view concerning permissiveness.

[137]See Kelly (2014), Meacham (2014), Schoenfield (2014) and Sharadin (2015). While Titelbaum endorses both intrapersonal and interpersonal permissiveness (see the previous footnote), he endorses diachronic requirements of epistemic rationality typically leading to the compromise between intrapersonal uniqueness and interpersonal permissiveness (Titelbaum 2013, chap. 7).

Table 4.2. Aggregated votes of epistemically rational agents on various propositions

	P	Q	(P ^ Q)
Member 1	Yes	No	No
Member 2	No	Yes	No
Member 3	Yes	Yes	Yes
Aggregation	Yes	Yes	No

In the above table, we can assume that every agent responded correctly to his or her epistemic reasons, but that they find themselves in an epistemically permissive situation. Nevertheless, the group agent ends up with an inconsistent combination of attitudes through aggregation (and a Closure principle). So, in interpersonal permissive situations, Consistency might not be a by-product of other requirements.[138]

Buchak and Pettit then analyze an objection from putative "epistemic restrictions" upon group agents. It could be argued that the suggested aggregation procedure is inconclusive, since a more fine-grained aggregation procedure would preserve consistency. For instance, agents could provide their degrees of belief in P and Q instead of voting "Yes" or "No". The aggregation procedure would lead to group degrees of belief in P and Q. However, Buchak and Pettit reject this possibility, stating that even epistemically rational agents cannot provide fine-grained credences in the propositions they believe. Such a solution would be unrealistic (Buchak and Pettit 2015, 221–22).

The upshot here is that, insofar as interpersonal permissiveness is true, Consistency might play an explanatory role in collective contexts such as voting groups. That is, rational agents who share the same epistemic reasons might not vote unanimously on various issues, and thus end up with inconsistent combinations of group beliefs (or collective decisions). In order to avoid such a result, we need a requirement prohibiting group agents (or agents under simple majority rule) from entertaining such inconsistent combinations of beliefs. But once

[138]Relatedly, Worsnip (m.s.) argues that disagreement can be interpreted as interpersonal incoherence. Specifically, two agents disagree on whether P if they hold beliefs (or other attitudes) that a single individual would be incoherent to entertain.

again, such a requirement can't necessarily be derived from a requirement of responding correctly to reasons one has, since agents can be in an interpersonal permissive situation.

Of course, as we can see in the above table, Buchak and Pettit accept a Rational Closure principle, according to which if one is rational in believing P and rational in believing Q, then one is rational in believing (P^Q). This is why the collective agent ends up with an inconsistent combination of beliefs. As I explained in the previous chapter, Rational Closure is problematic, especially in lottery cases. Buchak and Pettit are aware that such an assumption is contentious (Buchak and Pettit 2015, 212). But even if Rational Closure is false, the possibility of interpersonal permissiveness might lead epistemic groups to entertain "close to inconsistent" combinations of beliefs (such as believing P, believing Q and disbelieving (P^Q)). So minimally, if interpersonal permissiveness is true, a requirement that is similar to Consistency—that is, a requirement that prohibits "close to inconsistent" combinations of beliefs—might play an explanatory role in collective contexts.

4.4.3. Consistency Without Permissiveness? The Case of Collective Contexts

Now, could Consistency play an explanatory role in collective contexts even if Extreme Reasons Permissiveness is false? Buchak and Pettit think so. Indeed, they argue that there can be group agents where members of the group are not epistemic peers sharing all relevant epistemic reasons. This is so, because there can be group agents without assuming that there is a unique body of epistemic reasons shared by all the members of the group (Buchak and Pettit 2015, 222–24). Accordingly, one interpretation of table 4.2 is that different agents with different epistemic reasons vote differently (we can assume that their epistemic reasons support distinct incompatible conclusions). However, the aggregation of the votes is inconsistent (or close to inconsistent). In view of the foregoing, Buchak and Pettit conclude that, even if Extreme Reasons Permissiveness is false and the members of the group respond correctly to their epistemic reasons, Consistency could play an explanatory role in prohibiting the group from forming inconsistent combinations of beliefs.

I do not believe that the above argument is conclusive. Specifically, I believe that, if Extreme Reasons Permissiveness is false, Consistency doesn't play a distinct explanatory role in collective contexts. Here is why.

The above line of reasoning does not take into account the epistemic significance of the voting procedure itself. If Extreme Permissiveness is false, the fact that agents take distinct incompatible attitudes towards P can bear epistemic significance. Indeed, apart from rational mistakes, there are three explanations of why agents take incompatible attitudes towards P: (i) they do not have the same epistemic reasons, (ii) one of them is irrational or (iii) they find themselves in a permissive situation.[139] Now, assume that the members of a given group are rational, that Extreme Permissiveness is false, and that such facts are common knowledge. In such a case, when agents learn that they disagree concerning P, this means that at least one of them has epistemic reasons unavailable to the others. Presumably, if we learn that others have relevant information concerning P which is unavailable to us (through disagreement, deliberation or other means), this should affect our own attitude towards P. In short, learning that others disagree with us can affect our degree of confidence in P.

With respect to Buchak and Pettit's argument, why does the epistemic significance of disagreement matter? If Extreme Permissiveness is false, learning the results of a vote can affect which attitudes it is rational for agents to have. If agents *learn something* from their disagreement with each other, a good aggregation procedure should reflect such an acquisition of new information. Presumably, if members of a collective are anything like a group, they should care about the results of the vote. However, Buchak and Pettit's aggregation procedure does not take such a possibility into account. If P is supported by a very thin majority, then following the aggregation procedure, it will be assumed that the group agent accepts P. And once the vote has been taken, agents cannot change their vote even if the vote's results constitutes new information. This is problematic, since in realizing that P is only accepted by a thin majority, some members of the group might revise their attitude towards P. Provided that

[139]This is a common assumption in the literature surrounding peer disagreement. See, for instance, Christensen (2009), Kelly (2005) and Schoenfield (2014).

Extreme Permissiveness is false, the aggregation procedure should be dynamic, especially in cases where members of a group realize that P is only supported by a razor-thin majority.

In summary, I believe that Consistency plays an explanatory role in collective contexts only if interpersonal permissiveness (of the extreme sort) is true. However, if interpersonal permissiveness is false, it is implausible that Consistency plays an explanatory role in collective contexts. Once again, this supports the Modest Reasons-Responsiveness Thesis.

4.5. Conclusion: Some Progress Towards the Modest Reductionist Hypothesis

In this chapter, I argued that the Modest Reasons-Responsiveness Thesis is correct. According to such a view, Consistency plays a distinct explanatory role in cases where both believing P and disbelieving P are warranted by a body of epistemic reasons. I conceded that, relative to a body of epistemic reasons and some epistemic standards, permissiveness (of the extreme sort) is wildly implausible. Still, one possibility remains: there could be distinct but equally rational epistemic standards available to agents. Hence, Consistency might play an explanatory role in prohibiting inconsistent combinations of epistemic standards.

Now, does that mean that the Modest Reductionist Hypothesis is false? I think not. Recall that the Modest Reductionist Hypothesis is concerned with *ideal* theories of epistemic rationality. Perhaps non-ideal theories of epistemic rationality are permissive while ideal theories of epistemic rationality are impermissive.

In the next chapter, I will argue that the Permissive Epistemic Standards Thesis is false in ideal theories of rationality. Given what I have argued for in this chapter, this means that Consistency doesn't play an explanatory role in ideal theories of epistemic rationality. However, Consistency might very well play an explanatory role in non-ideal theories of epistemic rationality. Hence, the Modest Reductionist Hypothesis will be vindicated... but only in the next chapter.

Chapter 5. Against Permissive Epistemic Standards (for Ideal Agents)

Chapter summary. In this chapter, I argue that the Permissive Epistemic Standards Thesis (introduced and discussed in section 4.3.2) is false with respect to ideal theories of epistemic rationality. Specifically, those who argue that epistemic standards are permissive can't make sense of the reliability criterion (at least in ideal theories of epistemic rationality). My strategy relies on Condorcet's Jury Theorem. I then explain why my strategy might not generalize to non-ideal theories of epistemic rationality.

> Il ne s'agit pas ici de moi seul, mais de tous;
> je ne dois donc pas me conduire d'après ce que je crois être raisonnable,
> mais d'après ce que tous, en faisant comme moi, abstraction de leur opinion,
> doivent regarder comme étant conforme à la raison et à la vérité.
>
> Condorcet, *Essai* (cvii)

William, Harry and Melania just finished analyzing the evidence for a study in criminology. Melania believes that their evidence supports the conclusion that P, while Harry and William do not. Here is an explanation of why they disagree: they have the same relevant evidence, but they have different epistemic standards. The evidence is conclusive *relative to Melania's epistemic standards*, but the evidence is inconclusive *relative to William's or Harry's epistemic standards*.

If the Permissive Epistemic Standards Thesis is true, it is perfectly possible that William, Harry and Melania are fully rational. Indeed, according to such a view, the weight of the epistemic reasons in favour of P is not objectively determined—rather, in order to determine their weight, epistemic reasons have to be subjectively mediated through an agent's rational epistemic standards. Furthermore, there are multiple incompatible rational epistemic standards, in the sense that there are incompatible but equally rational ways to reason from a given body of epistemic reasons. Hence, agents like William, Harry and Melania can entertain rational disagreement on whether P, and such disagreements are explained by the fact that they entertain distinct incompatible rational standards.

In this chapter, I challenge the Permissive Epistemic Standards Thesis. In section 5.1, I discuss various unsatisfactory arguments against the Permissive Epistemic Standards Thesis. In section 5.2, I present my own argument against such a view. I begin by discussing the relationship between the Permissive Epistemic Standards Thesis and the reliability criterion, which roughly states that an ideally rational agent's epistemic standards optimize his or her ratio of true to false beliefs.[140] I then present Titelbaum and Kopec's Reasoning Room, which seems to support the claim that the Permissive Epistemic Standards Thesis is compatible with the reliability criterion. If cases like the Reasoning Room are conclusive, they provide support for the claim that there can be equally reliable but incompatible epistemic standards. In sections 5.2.2 and 5.2.3, I argue that Titelbaum and Kopec's Reasoning Room gives rise to an important objection against the Permissive Epistemic Standards Thesis. My strategy relies on Condorcet's Jury Theorem. Relying on this theorem, I will argue that, no matter how many incompatible epistemic standards are equally reliable, one epistemic standard is more reliable than all such standards. More generally, this means that the reliability criterion tends to support the denial of the Permissive Epistemic Standards Thesis, since ideally rational epistemic peers should endorse the same maximally reliable epistemic standards.

In section 5.3, I explain why the above argument might fail to generalize to non-ideal theories of epistemic rationality. That is, the argument works for *ideally* rational agents—that is, agents with great cognitive capacities who satisfy the optimal rational epistemic standards available to them. By way of contrast, many theories of epistemic rationality suggest that agents are permitted to satisfy suboptimal but sufficiently good epistemic standards, such as fast and frugal heuristics (Todd and Gigerenzer 2000). My argument does not apply to non-ideal theories. In section 5.4, I respond to an objection developed by Miriam Schoenfield and others. In section 5.5, I explain how my argument against the Permissive Epistemic Standards Thesis closes my case in favour of the normativity of epistemic rationality.

[140]As I briefly explain in section 2.2.1, the ideal of epistemic rationality does not necessarily match the absolute epistemic ideal. In a teleological perspective, the best belief-forming process would be something similar to: "A believes P if and only if P is true." By way of contrast, the ideal of epistemic rationality allows for false beliefs.

5.1. Unsatisfactory Arguments Against the Permissive Epistemic Standards Thesis

5.1.1. The Objection From Arbitrariness

Here is a first objection against the Permissive Epistemic Standards Thesis. Roger White (2005) argues that, if you think you are in an epistemically permissive situation, and that, between two incompatible attitudes or standards, you may choose either one without violating rationality requirements, then your choice will be arbitrary. There will be no significant difference between (a) your own choice and (b) popping a belief-inducing pill and reducing your belief-formation process to a random procedure (White 2005, 448). Since a random procedure is not sensitive to epistemic reasons, White argues that making a choice between different putative permitted options would violate rationality requirements. In other words, fully rational agents follow procedures that are reasons-sensitive, but in a putative permissive situation, a procedure that is not reasons-sensitive (a random one) could lead an agent to make a choice between different permitted options. Therefore, the agent would not count as fully rational if her belief-formation procedure is not reasons-sensitive.

White's argument is intuitively correct at first sight. Unfortunately, such an argument is inconclusive in the cases that I discussed in section 4.3.1. Indeed, suppose that Kolodny's decision-theoretic argument is correct, and that one ought to optimize expected epistemic value. Then necessarily, there are cases where an arbitrary procedure will be perfectly fine. Indeed, as it was made clear in section 4.3.1, the expected value of believing P can be equal to the value of not believing P.[141] In such cases, one will *run out of reasons* to make an optimal decision. Furthermore, suspending judgment on whether we should believe P or not believe P won't be an option. Necessarily, you will believe P or you will not believe P. Thus, in such a case, an agent is running out of non-arbitrary procedures, and so there is simply no other correct procedure than an arbitrary one.

[141]To be clear: this would be a permissive situation of the moderate sort. Still, the objection from arbitrariness is unsuccessful in such cases. So, it is hard to see why the objection would be successful in permissive situations of the extreme sort.

My point here is that arbitrariness could very well be part of a rational belief-forming process, especially in cases where an agent has no other choice than making an arbitrary decision. So, White's objection doesn't seem satisfactory after all.

5.1.2. The Objection from Planning

Here is another objection against permissiveness. According to Daniel Greco and Brian Hedden (2016), a central function of our epistemic practices is closely related to planning, the activity of making decisions about some possible future worlds. There could be permissive and impermissive planning, as in the following:

Impermissive Planning. if an agent concludes that believing P is rational relative to evidence E, then he or she plans to hold P when acquiring evidence E.

Permissive Planning. if an agent concludes that believing P is rational relative to evidence E, then he or she plans not to exclude holding P when acquiring evidence E.

What's wrong with Permissive Planning? According to Greco and Hedden, if we frequently find ourselves in epistemically permissive situations, there will be "strict limits on just how committal subjects' doxastic plans can be—we'll have to say that in permissive cases, subjects' plans cannot rule out any of the permissible attitudes" (Greco and Hedden 2016, 380). However, in permissive situations, agents are rational to choose between incompatible doxastic attitudes (including epistemic standards). Thus, permissivists seem ill-equipped to make sense of some functions of our epistemic practices, such as planning.

But once again, Greco and Hedden's argument does not seem conclusive in the case discussed in section 4.3.1. Indeed, suppose that one ought to optimize expected epistemic value, but that two doxastic options have equal expected epistemic value. If, from an agent's perspective, there really is no ground for concluding that one option is more rational than the other, *no planning seems useful or required*. Suppose once again that you are offered the mutually exclusive options A and B. Option A gives you $1 and option B also gives you $1. In such a case, do you really need to plan in advance what option you will take? Will having a plan make your situation better? Planning appears unnecessary or useless in such cases, especially since there is no substantial difference between these options. From your

perspective, indifference between options A and B is perfectly rational. So the planning argument seems inconclusive in the cases I identified.

Furthermore, Greco and Hedden's claim that "subjects' plans cannot rule out any of the permissible attitudes" does not support the conclusion that agents cannot choose between incompatible doxastic attitudes. One could believe P without excluding not believing that P. For example, an agent could believe that P but believe that he could later change his mind. Again, take a case from the practical realm. Suppose that Buridan's ass is planning what to intend when facing identical stacks of hay. Call these options Left and Right. Let's assume that Buridan's ass plans to go Left. Did Buridan's ass *exclude* Right because it intended to go Left? Not necessarily, since abandoning an intention is possible without violating a rationality requirement. As long as planning or forming an intention does not affect which options are optimal, revising one's attitudes is correct.

Hence, as with White's objection from arbitrariness, the Planning argument doesn't seem conclusive.

5.1.3. The Objection from the Practical Relevance of Epistemic Practices

Here is a third argument against permissiveness (and, specifically, the Permissive Epistemic Standards Thesis). Sinan Dogramaci and Sophie Horowitz (2016) as well as Daniel Greco and Brian Hedden (2016) considered the practical significance of coordinating our epistemic standards. In other words, instead of focusing exclusively on the nature of epistemic norms, they take the practical dimension of our epistemic practices as a starting point for denying permissiveness. I will here focus on the version of the argument offered by Dogramaci and Horowitz.

Dogramaci and Horowitz begin by presenting a fairly common datum, namely that "our social practice of epistemically evaluating one another's beliefs has value" (Dogramaci and Horowitz 2016, 131). Call this the "initial datum". So far, the initial datum is ambiguous: the notions of social practices, evaluations and value are unclear. Let's see how Dogramaci and Horowitz understand these notions.

First, social practices refer to interpersonal practices among the members of an epistemic community. Dogramaci and Horowitz think that as members of an epistemic community, we regularly argue and discuss with each other because it is valuable, as members of a group, to evaluate each other's doxastic attitudes (ibid., 132).

Now, the epistemic evaluations Dogramaci and Horowitz are interested in are promoting rational beliefs and discouraging (or criticizing) irrational beliefs (ibid., 131). Judging that a belief is irrational typically means that such a belief wasn't formed in accordance with the requirements of rationality. In view of the foregoing, Dogramaci and Horowitz think that, in promoting and criticizing each other's attitudes, we are in fact evaluating the *rules* licensing certain beliefs relative to a body of evidence (ibid.). We can understand these rules as *epistemic standards*. For example, suppose that an agent believes P every time he or she has a rational credence of more than 0.95 in P. This means that he or she follows an epistemic standard such as "if I judge that P's probability is greater than 0.95, then I believe P." If a community judges that such an epistemic standard is incorrect, they could criticize the agent for believing P, but what they are ultimately criticizing is the epistemic standard underlying the belief that P.

It should be noted that, according to Dogramaci and Horowitz, there is a strong connection[142] between reliability and rational epistemic standards (ibid., 135). A belief-formation process is reliable when following such a process makes it more likely that agents will end up with true beliefs. Since Dogramaci and Horowitz think that there is a strong connection between rational epistemic standards and reliable processes, they conclude that promoting rational epistemic standards "make it more likely that one another's beliefs will be true" (ibid., 135).

The last notion to clarify is the value associated with such epistemic practices. Dogramaci and Horowitz think that there is a *practical purpose* to promoting rational epistemic standards (ibid., 136). Reliable testimony has practical value, since epistemic

[142]Dogramaci and Horowitz argue that, while there is a strong connection between rational epistemic standards and reliable processes, reliability is not a sufficient condition for epistemic rationality (Dogramaci and Horowitz 2016, 135). In such a context, we are left with two explanations of the connection between the two. Either reliability is a *necessary* condition for rationality or there is a *correlation* between reliability and rationality.

communities are interested in getting significant truths (or truths about matters of interest). An efficient way of getting these significant truths is to divide the epistemic labour of collecting evidence, reasoning and drawing conclusions among members of an epistemic community (ibid., 136-7). For example, if I am rational to conclude that P and you are rational to conclude that P implies Q, we could share our respective conclusions to reach a new rational conclusion, namely Q. In view of the foregoing, if members of an epistemic community end up with reliable epistemic standards, they are more likely to reach conclusions that others can trust, which serves the group's practical aim of getting significant truths.

We can now disambiguate the intial datum, which stated that "our social practice of epistemically evaluating one another's beliefs has value" (ibid., 131). With a clearer picture of the notions of social practices, evaluations and value, we can now reformulate the datum in the following way:

Reformulated Datum. In an epistemic community, promoting and criticizing each other's attitudes with respect to their epistemic rationality makes it more likely that agents collect evidence, reason, or draw conclusions others can trust.

The question is now whether uniqueness best explains such a datum. If members of epistemic communities reason from a *unique* set of rational epistemic standards, they can treat each other as epistemic surrogates, namely as agents with sufficiently similar modes of reasoning. In view of the goal of getting significant truths, having epistemic surrogates is valuable, since it allows agents to efficiently "divide the labor of *collecting evidence* and the labor of *reasoning*." (ibid., 138). In other words, since epistemic surrogates have sufficiently similar reliable modes of reasoning, they can provide reliable information to each other through testimony.

However, Dogramaci and Horowitz argue that this is not the case with permissive epistemic standards. If agents reason from a *permissive set* of epistemic standards, they will not be able to treat each other as epistemic surrogates. Indeed, in such a case, rational agents could have distinct incompatible epistemic standards, and they would constantly be required to review each other's standards to reach a conclusion. Consequently, Dogramaci and Horowitz conclude that uniqueness best explains their datum:

> So our explanation works, given uniqueness. And in fact, it *requires* uniqueness. If permissivism is true, then rational reasoners need not conform. That is, there are cases where rational reasoners use alternative belief-forming rules, rules that yield distinct views given the same evidence. In this case, the enforcement of rational rules of reasoning does not make it safe to trust the testimony of rational reasoners, since there is now a risk that a rational reasoner will not be reliable. (ibid., 139)

In summary, the Epistemic Practices Argument seems to support uniqueness, in the following way:

(1) In an epistemic community, promoting and criticizing each other's attitudes with respect to their epistemic rationality makes it more likely that agents collect evidence, reason, or draw conclusions others can trust;

(2) The most efficient way to collect evidence, reason, or draw conclusions is for agents to treat each other as epistemic surrogates and to divide the epistemic labour among themselves;

(3) Agents' treatment of each other as epistemic surrogates and the division of the epistemic labour among themselves is best explained by uniqueness;

(C) By the principle of inference to the best explanation, it follows that uniqueness is true.

However, such an argument is implausible. What if a distinct epistemic practice does not require that we treat each other as epistemic surrogates and is best explained by permissiveness? This would lead to a dilemma: either our epistemic practices do not correctly reflect the requirements of epistemic rationality, or they do. Taking the first horn of the dilemma means that there is an explanatory gap in the Epistemic Practices Argument. Taking the second horn of the dilemma means that some types of epistemic labour are best explained by permissiveness, and so once again the Epistemic Practices Argument is compromised.

First, consider the first horn of the dilemma. If our epistemic practices do not reflect the requirements of epistemic rationality, we cannot determine the standards of epistemic rationality by observing our epistemic practice. This means that there is an explanatory gap in the Epistemic Practices Argument. Recall that, according to defenders of the Epistemic Practices Argument, the *fact* that we treat each other as epistemic surrogates is best explained by uniqueness. However, uniqueness is a *normative* thesis stating that epistemic peers concerning P are *rationally required* to come to the same conclusions on whether P. Deriving a

normative conclusion from mere factual considerations is problematic, since we cannot assume that what *is* the case is *valuable*, *required*, or *permitted*. To put it differently, we should not assume that, in observing our epistemic practices, we necessarily learn something about normative requirements.

Here is another way to put it. Defenders of permissiveness might accept that treating each other as epistemic surrogates is best explained by the belief in uniqueness. For instance, there could be general acceptance of uniqueness among the population, and this could explain why, as a matter of fact, we treat each other as epistemic surrogates. However, since defenders of permissiveness think that uniqueness is false, they will simply conclude that such an epistemic practice is irrational, mistaken or that it goes beyond what is rationally required of agents. Specifically, the Epistemic Practices Argument implicitly presupposes that communities are composed of epistemically rational reasoners and that interactions within communities *correctly reflect* rationality requirements. Without such a presupposition, inferring uniqueness from our epistemic practices is unjustified.

Let's now pass to the second horn of the dilemma by assuming that, in observing our epistemic practices, we can determine what is rationally required of agents. Now, the problem is that if we make such an assumption, we can find epistemic practices supporting permissiveness. Many epistemic communities encourage (or at least maintain) methodological diversity and heterogeneity. In philosophy of science, for example, there is a divide concerning the norms of universalism and pluralism, as Helen Longino explains in the following:

> Researchers committed to a monist or unified science will see plurality as a problem to be overcome, while researchers already committed to a deeply social view of science will see plurality as a resource of communities rather than a problem... Universalism and unification require the elimination of epistemologically relevant diversity, while a pluralist stance promotes it and the deeply social conception of knowledge that follows. (Longino 2016, sec. 4)

Not only in science do we encourage diversity and heterogeneity. All things being equal, we invite people holding different standards to a public debate, we praise dissenting philosophers

for diversifying the perspectives on a given question, we leave scientists free to use distinct methods leading them to incompatible conclusions, we ask lawyers to hold distinct standards regarding a litigation (a lawyer should always take his or her client's side and argue accordingly), and so forth. As epistemic communities, we sometimes promote incompatible epistemic standards. While we think that there is value in confronting them, such a confrontation doesn't always aim at reconciling everyone's standards. Confrontation may have a different goal, like making it salient that some standards of reasoning are incompatible. In such specific contexts, our epistemic practices suggest that epistemic heterogeneity bears value.[143]

Where does that leave us? In cases where we promote methodological heterogeneity, the division of some epistemic labour could make sense *without* uniqueness. For example, epistemic communities value critical thinking, and it is possible that a necessary mean of developing critical thinking is to confront incompatible epistemic standards with one another in the public sphere. Consider the case of public debates. In confronting different perspectives with each other, a public can realize that there are numerous distinct ways to reason on a given issue. So, in holding different epistemic standards and confronting them, debaters are *useful* to epistemic communities, since they help reinforce the public's critical skills. Furthermore, debating is a type of epistemic labour. For that reason, eliminating epistemic diversity within an epistemic community can result in blocking a fruitful type of epistemic labour.

Thus, there is a sense in which the division of epistemic labour is entirely compatible with the fact that agents are not epistemic surrogates. In cases like debating, holding distinct epistemic standards (and so, not treating each other as epistemic surrogates) is not an obstacle to accomplishing collective epistemic labour. In fact, a necessary condition for fruitful debates is that debaters do not defend exactly the same arguments or use the same methods (more on this point in section 5.3).

Defenders of the argument from the practical relevance of epistemic practices then face a dilemma. Indeed, at least one of the following is true: either (i) we cannot determine the

[143]See also Hong and Page (2012), Landemore (2012; 2013), Page (2007; 2010), Sunstein (2006) or Surowiecki (2005) on the epistemic benefits of diversity.

standards of epistemic rationality by observing our epistemic practices or (ii) we can (because our epistemic practices correctly reflect rationality requirements). Taking the first horn of the dilemma, one cannot draw normative conclusions concerning our epistemic practices from facts concerning our epistemic practices. However, the argument's third premise states that the *fact* that we treat each other as epistemic surrogates is best explained by a *normative* thesis like uniqueness. In view of the foregoing, the first horn of the dilemma leads to rejecting the third premise of the Epistemic Practices Argument. Taking the second horn of the dilemma, one can assume that we can draw normative conclusions from our actual epistemic practices. However, since some epistemic practices are better explained by permissiveness, this means that our epistemic practices do not support uniqueness (at least, since some epistemic practices are better explained by permissiveness, it should not be argued that our epistemic practices support uniqueness). Either way, the Epistemic Practices Argument is compromised.

5.2. A New Argument Against the Permissive Epistemic Standards Thesis

5.2.1. Reliability in the Reasoning Room

Many authors (including me) doubt that the Permissive Epistemic Standard Thesis is compatible with the reliability criterion. Before I present such an objection, I will give a rough account of the relationship between reliability and epistemic rationality. As I indicated in section 1.5.4, a necessary condition for ideal epistemic rationality is reliability—that is, I have endorsed the following:

Reliability Criterion. In the right conditions, if A is ideally rational, then A satisfies some available epistemic standards that optimize his or her ratio of true to false beliefs (and such standards lead A to reach the right answer more than 50% of the time).

Those who deny the Permissive Epistemic Standards Thesis can then offer the following objection: if there were incompatible but rational epistemic standards, they would not satisfy the Reliability Criterion. If I am permitted to believe P and you are permitted to disbelieve P relative to the same epistemic reasons, it seems that I don't have more than a 0.5 chance of

getting the right answer, and you don't have more than a 0.5 chance of getting the right answer.[144] If an agent who satisfies an epistemic standard has a 0.5 chance of getting the right answer, such an epistemic standard is unreliable. Hence, it could be argued that the Permissive Epistemic Standards Thesis is incompatible with the Reliability Criterion.

Titelbaum and Kopec (forthcoming; m.s.) have argued that the above line of reasoning is inconclusive. Indeed, imagine a group of epistemic peers who reason independently of each other from distinct incompatible epistemic standards. According to Titelbaum and Kopec, such agents can be equally reliable, in the sense that they can have an equal probability of being correct. The Reasoning Room case illustrates such as possibility:

Reasoning Room. "You are standing in a room with nine other people. Over time the group will be given a sequence of hypotheses to evaluate. Each person in the room currently possesses the same total evidence relevant to those hypotheses. But each person has a different method of reasoning about that evidence. When you are given a hypothesis, you will apply your methods to reason about it in light of your evidence, and your reasoning will suggest either that the evidence supports belief in the hypothesis, or that the evidence supports belief in its negation.... For each hypothesis, 9 people reach the same conclusion about which belief the evidence supports, while the remaining person concludes the opposite[145] .... [E]ach person in the room takes the evidence to support a belief that turns out to be true 90% of the time" (Titelbaum and Kopec forthcoming, 14).[146]

As we can see in the above case, the members of the group are equally reliable. Each of them reaches the right answer 90% of the time. Yet, with respect to each hypothesis

[144]See White (2005, 2014), Ballantyne (2012) and Dogramaci and Horowitz (2016) on this worry. See Titelbaum and Kopec (forthcoming, 12-19) for discussion.

[145]This aspect of Titelbaum and Kopec's argument could be clarified. Given the fact that agents reason independently from distinct incompatible standards, we should not expect exactly 9 agents to come to the same conclusion every single time. I leave this worry aside here. What matters in this chapter is that, *on average*, 90% of agents get the right answer.

[146]What about cases where the hypothesis considered is trivially true (for example, what if the hypothesis considered is "there is currently more than one person in the Reasoning Room")? In such cases, it seems everyone in the Reasoning Room should reach the same conclusion. Plausibly, what Titelbaum and Kopec have in mind is that the hypotheses considered in the Reasoning Room are not trivially true. Besides, their argument is conclusive insofar as it applies to *some* hypotheses, not *all* hypotheses. One could also wonder why agents in the Reasoning Room never withhold judgment concerning P. Indeed, in some situations, a rational agent's epistemic standards can recommend neither believing nor disbelieving P. According to Titelbaum and Kopec, such situations do not happen in the Reasoning Room, since one's reasoning "will suggest either that the evidence supports belief in the hypothesis, or that the evidence supports belief in its negation" (Titelbaum and Kopec forthcoming, 14). In order to accommodate the possibility of rational suspension of judgment, we can assume that agents who withhold judgment concerning a given hypothesis momentarily leave the Reasoning Room when such an hypothesis is evaluated.

presented to the participants, there is no consensus among them on which answer is right. So, this means that agents with different incompatible epistemic standards can be equally reliable. Therefore, in accordance with the Reasoning Room case, the Permissive Epistemic Standards Thesis seems compatible with the Reliability Criterion.

5.2.2 Introducing Condorcet's Jury Theorem

In the remainder of this section, I will argue that attempts to reconcile the Permissive Epistemic Standards Thesis and the Reliability Criterion (as in Titelbaum and Kopec's Reasoning Room) are subject to an important objection. Specifically, I will argue that, no matter how many distinct reliable epistemic standards there are, at least one epistemic standard is more reliable than such standards. This result is a direct consequence of Condorcet's Jury Theorem.

Condorcet's Jury Theorem roughly states that when equally reliable voters make a choice between two outcomes (such as "guilty" and "not guilty"), the majority rule can lead them to make collective decisions "with a competence that approaches 1 (infallibility) as either the size of the group or the individual competence goes up" (Estlund 1994, 131). This result is correct under a few assumptions. First, the probability that voter i gets the right answer (Pr_i) satisfies the following: $0.5<Pr_i<1$ and $Pr_1=Pr_2...=Pr_n$. Second, the probability that a voter gets the right answer is independent of the probability of any other voter getting the right answer (Condorcet 1785).

The Theorem has interesting implications concerning group competence. Here is an example. Suppose that three voters are trying to determine if Jones is guilty (and suppose that Jones is guilty). Each voter has 2/3 chance to get the right answer. However, the probability that the group will reach the conclusion that Jones is guilty under the majority rule is 20/27 (Estlund 1994, 136).[147] Since 20/27>18/27, this means that the group competence using majority rule outperforms the individual competence.

[147]Suppose the jurors are A, B and C. The probability that A, B and C will get the right answer is 8/27 (or 2/3·2/3·2/3). The probability that only A and B will get the right answer is 4/27 (or 2/3·2/3·1/3). Relatedly, the probability that only B and C will get the right answer is 4/27, and the probability that only A and C will get the right answer is 4/27. Since these possibilities are mutually exclusive, the probability that the group under simple majority will get the right answer is 20/27.

Condorcet's Jury Theorem has received a great deal of attention in political philosophy, where it is argued that such a theorem confirms the "wisdom of crowds", or the ability of voting groups to identify the correct decision through simple procedures such as voting (Landemore 2012; List and Goodin 2001; Surowiecki 2005; Vermeule 2012, 344–45). It also has found applications in artificial intelligence for justifying ensemble methods, which consist in combining multiple learning algorithms. Indeed, distinct but equally good learning algorithms can give rise to a better algorithm under an ensemble method (Polikar 2012; Rokach 2010).

However, a common objection against the Theorem is that it is hardly applicable to real-life situations. First, the independence condition is rarely satisfied. For example, opinion leaders frequently influence others in voting like them, which violates the independence condition. Second, the equal reliability condition is unrealistic. For instance, some agents are dogmatic, are not very good at making clever inferences or fail to respond correctly to their epistemic reasons, while others competently weight their epistemic reasons, reason correctly or draw clever conclusions. Such factors influence one's reliability. In such a context, it seems implausible that a jury will happen to be composed of equally reliable agents (some agents will be *unreliable* and others will be *too reliable*!). Third, the Jury Theorem is applicable only insofar as agents are confronted with a binary choice (such as "believe P" and "disbelieve P"). However, in many situations, we often have more than two options (such as "Bush," "Gore" or "Nader").

In view of the above complications, several authors have tried to relax the Theorem's conditions, so that it can be applicable to real-life problems (Bachrach et al. 2012; Dietrich and Spiekermann 2013; Fey 2003; Kaniovski 2010; List and Goodin 2001; Romeijn and Atkinson 2011; Stone 2015). However, with respect to the project of this chapter, it appears that there is no need to relax the original Theorem's conditions. As I will explain shortly, the unrealistic premises of the original Theorem seem to be satisfied in cases like the Reasoning Room. So, while these relaxed versions of the Theorem are more than relevant, the original theorem will do the trick here.

5.2.3. Reinterpreting the Reasoning Room as a Jury Case

We can reinterpret Titelbaum and Kopec's Reasoning Room as a jury case. Indeed, in Titelbaum and Kopec's thought experiment, all of the Jury Theorem's conditions are satisfied. First, all the agents are equally reliable: the probability that reasoner i gets the right answer is greater than 0.5, and it is equal to the probability that reasoner j gets the right answer (for $i \neq j$). Second, the agents in the Reasoning Room face a binary choice ("believe P" or "disbelieve P"). Third, the agents are reasoning independently of each other and their standards of reasoning are incompatible, which implies that the independence condition is satisfied: with respect to the hypotheses analyzed in the reasoning room, the probability that a reasoner gets the right answer is independent of the probability of any other reasoner getting the right answer. Hence, the Theorem's conditions are satisfied. So, Condorcet's Jury Theorem can be applied to Titelbaum and Kopec's thought experiment.

It could be argued that, in the Reasoning Room, the independence condition is merely satisfied in part. Indeed, since agents have the same evidence, this could lead them to reach similar conclusions about the various hypotheses examined. In such a context, one could argue that agents are not independent of each other, since their sharing the same evidence will lead them to form correlated beliefs. Minimally, it would be more rigorous to claim that, in the Reasoning Room, independence is conditional on the evidence.

I have two responses to this worry. First, in Titelbaum and Kopec's Reasoning Room, it is false that agents tend to reach similar conclusions about the various hypotheses examined: for each hypothesis considered, there is no correlation among the answers. That is, the fact that agent i got the right answer does not make it more probable that agent j will get the right answer. So, even if agents share the same evidence, this does not entail that the independence condition is violated. Second, using a relaxed independence condition that is compatible with the fact that voters have the same evidence (such as Dietrich and List's (2004, 182) Independence Given the Evidence condition[148] or Dietrich and Spiekermann's (2013, sect. 4)

[148]Independence Given the Evidence states: "The votes V_1, V_2, ..., V_n are independent from each other, conditional on the body of evidence E" (Dietrich and List 2004, 182). Such a condition affects Condorcet's Jury Theorem in cases where the evidence is misleading. In other words, "the probability that the majority verdict matches the true state of the world (given that state) converges to the probability that the ideal interpretation of

New Independence condition[149]) would not affect the results of this chapter. These relaxed conditions aim at accommodating the possibility of facing very difficult problems or of having very misleading evidence, which influence whether agents will get the right answer. However, in Titelbaum and Kopec's Reasoning Room, the evidence is not abnormally misleading and agents do not face very difficult problems, since 90% of the participants get the right answer each time.[150] Given that such relaxed independence conditions do not affect the results of this chapter, for the sake of *simplicity*, I will assume that the original independence condition is satisfied in the Reasoning Room.

In accordance with such a reinterpretation of Titelbaum and Kopec's Reasoning Room, I will now argue that the following thesis is correct:

Superior Standard Thesis. Let $\{St_1, St_2, ..., St_n\}$ be a set of incompatible available epistemic standards containing n elements. Let Pr_i be the probability that satisfying St_i will lead one to form a true belief. Finally, assume that such standards are equally reliable (such that $Pr_1=Pr_2=...Pr_n$ and $0.5<Pr_i<1$). Then, there exists an available epistemic standard that is more reliable than St_1, St_2, ..., and St_n.

The argument is simple. Suppose that, in the Reasoning Room, Agent 1 satisfies St_1, Agent 2 satisfies St_2, and so forth. Then, we can design a "Condorcetian" epistemic standard, as in the following way:

(1) If a majority of agents in the Reasoning Room believe P, you should believe P.

(2) If a majority of agents in the Reasoning Room disbelieve P, you should disbelieve P.

Following the Jury Theorem, such an epistemic standard will necessarily be more reliable than St-1, St-2, ..., St-n. Indeed, in accordance with the Theorem, if a majority of agents believe that P, the probability that P is true is higher than Pr-i. Since $Pr_1=Pr_2=...Pr_n$, the Condorcetian standard is more reliable than any of the standards St_1, St_2, ..., St_n. Similarly, if a majority of agents believe that ~P, the probability that P is false is higher than Pr_i. Again, since

the evidence is correct, i.e., that the evidence is not misleading" (Dietrich and List 2004, 187).

[149]New Independence states: "The events R_1, R_2, ... that voters 1, 2, ... vote correctly are independent conditional on the problem π." Such a condition affects Condorcet's Jury Theorem in cases where the problem is very difficult, since agents then have less than 0.5 chance of solving the problem correctly.

[150]In any case, it is not even clear that defenders of the Permissive Epistemic Standards Thesis would accept the existence of such a thing as "misleading evidence", since they do not think that the evidence points in a direction on its own—the evidence points in a direction or another insofar as it has been mediated through an agent's epistemic standards.

$Pr_1=Pr_2=...Pr_n$, the Condorcetian standard would be more reliable than the standards St_1, St_2, ..., St_n. This means that, in any case, the Condorcetian standard outperforms the standards St_1, St_2, ..., St_n. Therefore, the Superior Standard Thesis is true as a direct consequence of the Jury Theorem.

The idea behind the Superior Standard Thesis is that distinct but reliable standards can be reconciled under a unique (and more reliable) "meta-standard." This line of reasoning is largely inspired by recent work in artificial intelligence, where the simple majority rule can be used to improve accuracy. Indeed, in order to answer a given question, some distinct but equally reliable algorithms (or classifiers) analyze the available data. When the algorithms do not reach consensus, an ensemble method based on majority voting can be used to reach a final verdict (Rokach 2010, sec. 3.1). For example, in a situation where all of the Jury Theorem's conditions are satisfied, if 10 out of 15 reliable algorithms reached the conclusion that P, the final verdict given by the ensemble method would be that P.[151] Following Condorcet's Jury Theorem, such a method outperforms individual algorithms, in the sense that it is more reliable than every individual algorithm. As Polikar (2012, 1–2) notes, whether we are concerned with a community's democratic choices or a robot's learning process, Condorcet's Jury Theorem can be mobilized in favour of the simple majority rule.

5.2.4. Conceivable and Available Standards

Following the Superior Standard Thesis, if the members of the Reasoning Room come to know the vote's result, it would be more reliable for them to go with the majority. But what about a case in which such results are unknown? Indeed, in Titelbaum and Kopec's version of the Reasoning Room, "it's unpredictable who will be the odd person out for any given hypothesis. The identity of the outlier jumps around the room" (Titelbaum and Kopec forthcoming, 14). Presumably, what they have in mind is that agents do not share their conclusions with each other. In such a case, agents might lack the capacity to evaluate their epistemic reasons with the Condorcetian epistemic standard. Since epistemic standards are the

[151]As I explained in section 5.2.2, it is rarely the case that all of the Jury Theorem's conditions are satisfied. In artificial intelligence, some algorithms are more reliable than others. Accordingly, the ensemble method is often more complex than a mere application of the majority rule.

kind of things agents use to reason and reach new conclusions, it must be possible for agents to evaluate their epistemic reasons with such epistemic standards.

In response to this worry, recall that epistemic standards act as functions mapping an agent's evidence onto doxastic attitudes towards P. So, in order to evaluate the evidence with the Condorcetian epistemic standard, agents in the Reasoning Room do not necessarily need to share their conclusions with each other. Agents merely need to determine if a majority of available epistemic standards $\{St_1, St_2, ..., St_n\}$ support the conclusion that P. Specifically, they can process their evidence through epistemic standards $\{St_1, St_2, ..., St_n\}$, see which doxastic attitude towards P is supported by a majority of epistemic standards and go with the majority.

It could be objected that processing the evidence through epistemic standards $\{St_1, St_2, ..., St_n\}$ is too complicated. Such a worry is not relevant here. To be clear: I am not committed to the view that the Condorcetian standard is a *realistic* way for agents to reason from a body of evidence (especially if we are concerned with agents with limited cognitive capacities). My point is merely that such a standard is *conceivable*: we can imagine how a single agent or a group of peers could satisfy such a standard. The fact that real-life agents do not satisfy such a standard does not indicate that such a standard does not exist—rather, this indicates that real-life agents fail to satisfy a conceivable and highly reliable epistemic standard.

It could then objected that, while the Condorcetian standard is *conceivable*, it is not necessarily *available* to agents. There are many reliable but unavailable standards, such as "believe P if and only if P." Such a standard is unavailable because agents do not know how to satisfy it. One is not rationally required to entertain standards that one is not in a position to know, precisely because such standards are unavailable to one.

First, recall that the Condorcetian standard supervenes on standards $\{St_1, St_2, ..., St_n\}$ entertained by the agents in the Reasoning Room. So, if the standards $\{St_1, St_2, ..., St_n\}$ are available to all the agents in the Reasoning Room, the Condorcetian standard should also be available. They just need to determine if a majority of available epistemic standards $\{St_1, St_2, ..., St_n\}$ support the conclusion that P (which they know how to determine).

Now, could it be possible that some agents in the Reasoning Room are not in a position to know all the standards $\{St_1, St_2, ..., St_n\}$? This depends on whether agents need specific information, training or education to acquire those standards.

There is a sense in which different agents have access to different standards. Agents learn all sorts of methods in the course of their lives. To name a few, here are two examples. There are different techniques for solving a Rubik's cube. Some players are only familiar with the Roux method, while others are only familiar the Fridrich method. This is so, because they went through different training. The same goes for numerical analysts, who are familiar with different numerical methods for processing the data. Analysts take different courses, read different books, and so they end up knowing different numerical methods. Similarly, it is plausible to assume that some epistemic standards can be acquired in the course of an agent's epistemic life. Information, training or education can explain why some of us know how to process the evidence in a specific way, while others are not familiar with this way of processing the evidence.

Yet, recall that agents in the Reasoning Room are epistemic peers who share all relevant epistemic reasons. So while agents with different epistemic reasons can't necessarily think of the same standards (or are not in a position to know the same standards), agents in the Reasoning Room are in a position to know the same standards. This is so, because they share the same relevant epistemic reasons, including the evidence relevant to knowing some epistemic standards.

In view of the foregoing, the Condorcetian standard is conceivable and available to agents in the Reasoning Room. When epistemic peers entertain incompatible but equally reliable epistemic standards, they are in a position to identify a Condorcetian standard that is more reliable than their own standards.

5.2.5. A Dilemma for Defenders of the Permissive Epistemic Standards Thesis

The Condorcetian epistemic standard leads to a dilemma for defenders of the Permissive Epistemic Standards Thesis. The core of the dilemma is this: in the right conditions (for instance, when agents are not deceived brains in vats), either the rational incompatible

epistemic standards described by permissivists satisfy the reliability criterion, or they do not. Taking the first horn of the dilemma means that rational incompatible epistemic standards collapse into a unique Condorcetian epistemic standard, which supports the denial of the Permissive Epistemic Standards Thesis. Taking the second horn of the dilemma means that permissivists concerning epistemic standards can't make sense of the reliability criterion, which makes such a type of permissiveness less plausible. Either way, the argument from permissive epistemic standards is compromised.

Consider the first horn of the dilemma. Suppose that rational epistemic standards are maximally reliable. Then, the Condorcetian epistemic standard contradicts permissiveness. Even if there are equally reliable but incompatible epistemic standards, we can use such epistemic standards to design an even more reliable epistemic standard. The conceivability of a Condorcetian epistemic standard confirms that verdict. In other words, if (i) rational epistemic standards have something to do with their reliability and (ii) there exists a unique maximally reliable epistemic standard, this entails that permissiveness concerning epistemic standards is false.

Here is another way to put it. Defenders of the Permissive Epistemic Standards Thesis think that there are reliable but incompatible epistemic standards. Such standards are supposed to be equally reliable, and so we can't make a difference between these standards in terms of their reliability. However, this can't support the Permissive Epistemic Standards Thesis: if there are incompatible but equally reliable epistemic standards, *there has to be an even more reliable epistemic standard*. So, there has to be one maximally reliable epistemic standard, which compromises the view that epistemic standards are permissive.

Let's now turn to the second horn of the dilemma by assuming that rational epistemic standards are not necessarily maximally reliable. On this assumption, it seems that a worry raised in section 5.2.1 is correct: endorsing the Permissive Epistemic Standards Thesis amounts to separating epistemic rationality and reliability. This is problematic. Without reliability, it seems that we are left with very minimal constraints on rational epistemic standards, such as the internal consistency of rational standards. However, many consistent

sets of standards do not seem rational. If a body of epistemic standards is rational insofar as it satisfies consistency, any skeptic, grue-projector or conspiracy theorist can be regarded as rational (as long as his or her epistemic standards are consistent), and there is something wrong with such a conclusion (Horowitz 2014b, 45).

Defenders of the Permissive Epistemic Standards Thesis then face a dilemma. Taking the first horn of the dilemma, rational epistemic standards are maximally reliable. However, even if several epistemic standards or standards of reasoning are equally reliable, there exists an epistemic standard that is even more reliable, which contradicts the Permissive Epistemic Standards Thesis. Taking the second horn of the dilemma, rational epistemic standards are not maximally reliable. However, under such an assumption, we lack plausible constraints on what counts as rational epistemic standards. For instance, if permissive epistemic standards amount to consistent epistemic standards, many unreasonable but consistent epistemic standards will count as rational. Either way, the Permissive Epistemic Standards Thesis is compromised.

5.3. The Superior Standard Thesis Might Not Generalize to Non-Ideal Theories

5.3.1. First Objection: Epistemic Supererogation and the Reliability Criterion

The above argument might fail to generalize to non-ideal theories of epistemic rationality. First, it could be argued that, while the Condorcetian epistemic standard is more reliable than the putative permissive epistemic standards described by Titelbaum and Kopec, this does not entail that non-ideal rational agents are required to adopt the Condorcetian standard. Indeed, we could introduce a distinction between the required and the supererogatory, as in the following:

Epistemic Supererogation. While some epistemic standards are more reliable than others, satisfying them is supererogatory. Agents are *praiseworthy* for satisfying such standards, but they are not *required* to. By way of contrast, satisfying some epistemic standards that are reliable to a certain degree is required (such as the standards described by Titelbaum and Kopec).

Epistemic Supererogation ends up weakening the Reliability Criterion: an agent can satisfy suboptimal epistemic standards while being sufficiently reliable. This strategy has been pursued by Li, who claims that we can imagine a "special type of agent who performs special epistemic acts—acts that involve levels of insight, intelligence, and imagination that even very rational agents can fail to achieve. But we aren't required to exhibit such epistemic virtues" (Li 2017, 2). According to him, part of the debate surrounding permissiveness has to do with a confusion between the supererogatory and the required.

Can we save the Permissive Epistemic Standards Thesis by making a distinction between the required and the supererogatory? Perhaps that, in order to be rational, agents with limited cognitive capacities are merely required to satisfy some "reliable enough" (though not maximally reliable) epistemic standards. In other words, perhaps the Condorcetian epistemic standard is supererogatory for non-ideal agents.

5.3.2. Second Objection: Second-Best Epistemology and Imperfect Epistemic Practices

The second reason why the argument might not generalize to non-ideal theories of epistemic rationality echoes an argument I presented in chapter 2: we might be facing a problem of second-best in which approximating norms for ideal agents is suboptimal. Here is why.

Ideal agents who share the same epistemic reasons can always improve their respective epistemic standards. By *combining their standards*, agents can give rise to a better standard. However, combining epistemic standards might require great cognitive capacities or very specific knowledge that non-ideal agents can lack. So, what should non-ideal agents do?

Perhaps non-ideal agents can entertain distinct incompatible standards, and let collective institutions or social mechanisms do the combinatory work for them. Epistemic communities will then have successful mechanisms for acquiring true beliefs. However, these mechanisms will be very different from the ones found in ideal epistemic worlds. Such mechanisms will be adapted to the capacities of non-ideal agents. Furthermore, while agents in

ideal worlds necessarily entertain the same epistemic standards, agents in non-ideal worlds should rather entertain distinct incompatible epistemic standards. So, the truth of the Permissive Epistemic Standards Thesis might depend on whether we are concerned with ideal or non-ideal epistemic rationality.

Here is an example of how this might work. Suppose there are three equally reliable standards available to three agents with the same epistemic reasons. Suppose furthermore that it would be too demanding to ask each agent to reason from the Condorcetian standard (e.g., the standard that results from combining the three aforementioned standards). Nevertheless, one way these agents can compensate for their imperfections is this: they can adopt one standard each and vote on various issues from time to time. This way, assuming they go along with the majority view, they will replicate the Condorcetian standard in a way that is not too demanding. However, this way of replicating the Condorcetian standard *supports permissiveness*, in the sense that it is effective insofar as agents entertain distinct incompatible epistemic standards.

Hence, there is good reason to think that my argument does not generalize to non-ideal theories of epistemic rationality.

5.3.3. A Second-Best Problem: The Paradox of the Erratic Juror

The above line of reasoning in favour of permissiveness in non-ideal worlds might give rise to dilemmas for non-ideal agents. Indeed, suppose agents entertain distinct incompatible standards and trust institutions to do optimal combinatory work for them. Then, some problems of optimization can arise in such communities. Indeed, some individuals in the group might face a dilemma between improving individual reliability and improving collective reliability. This is the paradox of the erratic juror.[152]

Here is how group reliability can be understood: assume that 'group beliefs' supervene on individual beliefs through aggregation functions (such as the simple majority rule). Then, groups are reliable if and only if they satisfy some available belief-forming processes which

[152]I discuss this paradox in Daoust (m.s. b).

optimize the ratio of true to false group beliefs (and such processes lead the group to reach the right answer more than 50% of the time).

The paradox is that the above account of group reliability can conflict with the Reliability Criterion. An initial, intuitive description of the paradox goes as follows. A judge is in charge of forming a jury for several trials. Given the evidence the jurors acquire during the trials, they ought to determine if the defendants are guilty. The trials are subject to three important rules. First, all jurors are presented with the same evidence and they ought to vote on the basis of the shared evidence only. Second, the jurors do not discuss with each other before casting their vote. Third, the jurors are faced with a binary choice (such as "Guilty" or "Not Guilty").

The judge picks two jurors with distinct but equally commonsensical standards of reasoning—call them William and Harry. While they do not have the same methods of reasoning, William and Harry reach the right answer 60% of the time. So, they are fairly reliable. The judge also picks Melania, an "Erratic Juror" with unorthodox (but often misplaced) standards of reasoning. Melania is less reliable than the other jurors—she reaches the right answer 40% of the time. The following table reflects the general voting pattern of the jurors:

Table 5.1. Verdicts

Juror/Case	**P_1**	**P_2**	**P_3**	**P_4**	**P_5**
Juror 1 (William):	Not guilty	Guilty	Guilty	Guilty	Not guilty
Juror 2 (Harry):	Not guilty	Guilty	Guilty	Not guilty	Guilty
Erratic Juror (Melania):	Not guilty	Not guilty	Not guilty	Guilty	Guilty
Verdict:	Not guilty	Guilty	Guilty	Guilty	Guilty

Let's assume that the right verdict in each case is "Guilty". Let's also assume that, on the long term, the vote pattern of the jurors is always similar to the one in the above table. As we can see, the group reaches the right answer 80% of the time, which is a good success rate. Also,

Table 5.1 confirms that jurors 1 and 2 reach the right answer 60% of the time, and that the Erratic Juror reaches the right answer 40% of the time.

After a series of verdicts, the Erratic Juror is informed that her odd methods of reasoning are less reliable than the ones entertained by Juror 1 or Juror 2. Even worse, she is informed that her methods of reasoning leads her to the right answer less than 50% of the time. As a convinced reliabilist, the Erratic Juror is unsatisfied. One day, she discovers the methods of reasoning entertained by the Juror 1. In accordance with the Individual Reliability Criterion, she decides to improve her individual reliability by changing her methods of reasoning. In other words, she adopts the methods of reasoning entertained by Juror 1.

However, this improvement in individual reliability affects the group's voting pattern. Indeed, after the Erratic Juror changed her methods of reasoning, she now votes like Juror 1, as in the following:

Table 5.2. Verdicts after the Erratic Juror changed her methods of reasoning

Juror/Case	P_1	P_2	P_3	P_4	P_5
Juror 1:	Not guilty	Guilty	Guilty	Guilty	Not guilty
Juror 2:	Not guilty	Guilty	Guilty	Not guilty	Guilty
Erratic Juror:	Not guilty	Guilty	Guilty	Guilty	Not guilty
Verdict:	Not guilty	Guilty	Guilty	Guilty	Not guilty

As we can see, the group reliability is down to 60% after the Erratic Juror changed her methods of reasoning. Of course, in accordance with the Individual Reliability Criterion, the Erratic Juror's reliability has improved. However, the group is now less reliable, which violates the Group Reliability Criterion. Hence, it appears that, in some situations, it is impossible to satisfy the Individual Reliability Criterion and the Group Reliability Criterion simultaneously. Apparently, optimizing group reliability may sometimes reduce individual reliability.

Here is why the above case is paradoxical. It seems that the Erratic Juror faces the following unsolvable dilemma: improve individual reliability or improve collective reliability.

In ethics, we are familiar with trolley problems where an agent can sacrifice herself in order to save five other agents.[153] Similarly, there are well-known cases in which someone has to be a villain so that the group remains moral.[154] At first sight, the Paradox of the Erratic Juror is structurally similar to such ethical conundrums: in order to improve the group's reliability, an agent could sometimes be required to "sacrifice herself" or "be the villain", and thus adopt standards that are unreliable at the individual level. But this doesn't seem right.

As we can see, the paradox may arise in non-ideal scenarios, where rational agents do not entertain the same standards and let institutions do the aggregative work for them (via, for instance, simple majority rule). Elsewhere, I discuss this paradox at length.[155] With respect to the project of this chapter, all I want to stress is that the problems encountered by non-ideal agents can be very different from the problems encountered by ideal agents.

5.4. Diachronic Prohibition, Immodesty and Coherence

Suppose Kate entertains standards X, but that she is told by a reliable informant that some standards Y and Z (that are incompatible with X) are equally reliable. Suppose that such standards are, in fact, as reliable as X. Given the Superior Standard Thesis, she now has access to a better standard. Still, is she epistemically permitted to change her own standards? Some philosophers deny this. This is so, because of a putative diachronic norm prohibiting agents from changing their epistemic standards, as in the following:

Diachronic Prohibition. If agent A has a specific set of rational standards at time t_0 and does not acquire new epistemic reasons between t_0 and t_1, A should refrain from changing his or her rational standards at time t_1.

If Diachronic Prohibition is true, defenders of the Permissive Epistemic Standards Thesis can resist the argument I put forth in section 5.2. Indeed, even if there exists a uniquely superior standard, rational agents should refrain from changing their own standards (even if their standards are less than maximally reliable). Hence, Diachronic Prohibition supports the

[153]According to one variant of the problem, one can deviate the trolley in one's direction in order to save five workers.
[154]Think of torture cases: there are situations in which, if one commits torture, one can improve the group's morality by preventing various collective tragedies.
[155]See Daoust (m.s. b).

Permissive Epistemic Standards Thesis: depending on the epistemic standards rational agents adopt in the course of their epistemic lives, they might be required to entertain distinct incompatible standards.

In this section, I argue that Diachronic Prohibition is implausible. For the sake of the argument, I won't assume that agents have access to a more reliable epistemic standard. I will merely assume that agents have access to distinct but equally reliable standards. Even in such a context, I will argue that rational agents should be permitted to change their own epistemic standards, which is incompatible with Diachronic Prohibition.

5.4.1. Diachronic Prohibition Is Implausible

I will start by offering two thought experiments (Robot Acquisition and Transplant). These thought experiments have the same relevant normative features. I will then argue that, if Diachronic Prohibition is true, we need to make a normative distinction between these cases, which is problematic.[156] This will lead me to reject Diachronic Prohibition.

To begin with, consider the following two cases:

Robot Acquisition. Kate has designed Alpha and Beta, two robots based on replicas of her brain system. She has implanted circuits and chips in her brain so that all the evidence and sensory experiences she gathers are directly transmitted to her robots. While Kate has a risk-neutral prior function, her robots respectively have a risk-averse prior function and a risk-seeking prior function. When a big company asked her which robot they should buy, Kate said that she followed the principles of permissiveness and that there is no uniquely optimal risk-based prior function (we can assume that Kate is unaware of the Condorcetian standard).[157] While they function differently, both robots satisfy the requirements of epistemic rationality.

Transplant. Kate discovers a credence transplant procedure. Specifically, she identifies a method by which she can replace her credence function with another one. Alpha and Beta, her robots, are perfect matches for a credence transplant, since they are based on exact replicas of her brain system and have updated their credences on an identical body of

[156]Christensen (1994; 2000) reaches similar conclusions.

[157]For instance, suppose the Brier score is the right inaccuracy measure. Suppose that P, Q and R are true. Then, a robot with the conservative credences Cr(P)=0.6, Cr(Q)=0.6 and Cr(R)=0.3 gets a score of 0.81. A robot with the liberal credences Cr(P)=0.8, Cr(Q)=0.8 and Cr(R)=0.145 gets a score of ≈0.81. Their score are roughly similar. The difference between the two robots is that one of them entertains credences that are more extreme (e.g., closer to 0 or 1) and thus is risk-seeking.

evidence. So, Kate intends to get a credence transplant, and she could get it from Alpha or from Beta.

The decisions involved in the above cases are the following: in Robot Acquisition, a big company could buy Alpha or buy Beta, and, in Transplant, Kate could get a credence transplant from Alpha or from Beta. I here assume that we cannot find significant normative differences between cases or decisions if there are no relevant factual differences between them. If the decisions involved in Robot Acquisition and Transplant rest on the same relevant considerations, it cannot be the case that the same decision is rational in Robot Acquisition while it is irrational in Transplant. Otherwise, we would be committed to a form of bootstrapping, a process by which reasons or obligations appear out of nowhere. I reject bootstrapping (at least, the putative reasons or obligations one gets from bootstrapping are not epistemically rational).

Robot Acquisition and Transplant have the same relevant features. Of course, the scenarios are a little different: a big company could *buy* one of the robots while Kate could *exchange* her credence function for one of the robots' credence functions. However, if there are pros and cons related to choosing one option over the other in the Robot Acquisition case, then the same pros and cons related to choosing one option over the other will obtain in the Transplant case. For example, are the robots consistent? Are they reliable or accurate? Do they reason well? Do they lose information over time? If these factors are relevant in Robot Acquisition, they are also relevant in Transplant. In short, from an epistemically normative point of view, the kind of decision the big company has to make is no different from the kind of decision Kate could make. In view of the foregoing, we should not make significant normative distinctions between these cases.

Transplant is a good case for determining if there are diachronic norms prohibiting an agent from changing his or her credence function over time. If there are such norms, Kate is rationally prohibited from going for the credence-transplant procedure, since she would be prohibited from changing her prior function. However, as we can see in Robot Acquisition, there is no uniquely optimal risk-based prior function. Assuming that interpersonal permissiveness is true, a big company would not make a suboptimal decision in buying Alpha

rather than Beta and vice versa. Since Alpha and Beta are based on rational systems, it is hard to see why Kate is prohibited from abandoning her credence function and going for one of theirs. After all, if it is just a matter of risk profile and Kate feels like going for a risky epistemic life, she should be permitted to adopt Beta's credence function.[158] Furthermore, since Alpha and Beta update their credences on Kate's body of evidence, Kate has no reason to think that changing her credence function would result in her losing information. Thus, assuming that it is equally optimal for a big company to buy Alpha or Beta, Kate is permitted to change her credence function.

In summary, in cases like Transplant, we lack an explanation of why Kate would violate a requirement of rationality if she adopted a different credence function. It seems implausible that there would be diachronic norms of epistemic rationality prohibiting an agent from changing his or her attitudes over time. So, at least in permissive situations, Diachronic Prohibition is implausible.

5.4.2. The Objection From Immodesty

Here is an objection against my argument. In Robot Acquisition, Kate thinks that distinct incompatible epistemic standards are equally optimal. One could reply that agents like Kate ought to believe that their own standards are more accurate or truth-conducive than the others. According to Schoenfield, when an agent like Kate adopts or entertains a set of epistemic standards, she would refuse to adopt other incompatible standards. Indeed, compared with other epistemic standards, hers would now *appear* to be more truth-conducive, to maximize accuracy or to minimize inaccuracy. Also, a change in epistemic standards over time will strike an agent as irrational because "although she knows that, later, she will not be violating her own standards (since she will have new standards), she does not now think that her later standards will be as likely to lead her to a true belief as her current ones" (Schoenfield 2014, 201). So, Kate should not think that Alpha and Beta entertain optimal credence functions: she should believe that her credence function is optimal, and that Alpha's and Beta's credence functions are suboptimal.

[158]See, for instance, Pettigrew (2016b) on the relationship between risk-based decisions rules and prior functions.

If this is correct, we have an explanation of why Diachronic Prohibition obtains: entertaining epistemic standards changes our perception of other standards. Specifically, entertaining epistemic standards leads epistemically rational agents to believe that such standards are more truth-conducive than others. Schoenfield's argument echoes the Immodesty condition: an immodest agent estimates that his or her beliefs and epistemic standards are the most accurate ones (relative to a body of evidence).[159]

Even if it is a fact that there are distinct incompatible but equally reliable epistemic standards, an agent could *rationally (but falsely) believe* that his or her standards are epistemically superior. Insofar as there are rational false beliefs, Immodesty can explain why rational agents falsely believe that their own standards are more truth-conducive.

Be that as it may, Schoenfield's argument is problematic for three reasons. First, the argument put forth by Schoenfield relies on a problematic understanding of Immodesty. Following David Lewis, Schoenfield understands Immodesty as follows:

Immodesty (initial). Relative to their evidence, agents take their own standards to be the most truth-conducive ones.[160]

Yet, I believe Immodesty should be understood as follows:

Immodesty. Relative to their evidence, agents take their own standards to be *among* the most truth-conducive ones.

Against David Lewis, I here assume that immodest agents take their standards to be *among the most* truth-conducive ones (and not to be *the most* truth-conducive ones). Here is why. Lewis offers the following argument against modesty:

> Suppose you did trust some non-immodest method. By definition, it estimates some competing method to be more accurate than itself. So if you really did trust your original method, you should take its advice and transfer your trust to one of the competing methods it recommends. It is as if *Consumer Bulletin* were to advise you that *Consumer Reports* was a best buy whereas *Consumer Bulletin* itself was not acceptable; you could not possibly trust *Consumer Bulletin* completely thereafter. (Lewis 1971, 56)

[159]See Elga (2010), Lewis (1971), Christensen (2013), and Mayo-Wilson and Wheeler (2016) on Modesty, Immodesty and Strict Immodesty.
[160]As in Lewis (1971, 55) and Schoenfield (2014, 201–2), notably.

According to Lewis, agents who entertain modest standards estimate that other standards are more accurate. However, this doesn't entail that immodest agents take their standards to be the most accurate ones. Immodest agents can estimate that there is more than one optimal set of standards. In other words, there could be distinct but equally optimal standards available to agents. So, the argument put forth by Lewis merely entails that agents should take their standards to be *part of the best* ones available. [161]

Compare: Buridan's ass is facing two equally good stacks of hay (and he knows they are equally good). Standard Left says "choose the stack of hay on the left," while standard Right says "choose the stack of hay on the right." Suppose he is practically immodest—say, he thinks that the decisions resulting from his practical standards are among the best ones available. Yet, his standards do not necessarily preclude the possibility that other standards are equally optimal. Buridan's ass can entertain standard Left while thinking that standard Right is also optimal. In fact, given what he knows, Buridan's ass is fully aware that both standards are equally optimal. Thus, Immodesty should be neutral on whether optimal options are uniquely determined or not.

If immodest agents merely assume that their standards are among the most truth-conducive ones, Schoenfield's argument fails. Indeed, William can very well believe that his standards are among the most truth-conducive ones. Yet if he also believes that Melania's standards are equally reliable, he is not required to maintain his own standards (he could very well adopt Melania's standards). Thus, Diachronic Prohibition doesn't succeed if we interpret the Immodesty criterion correctly.

Second, with respect to acquired epistemic standards, a change in perception of rational standards leads to puzzling situations. Here is why.

It is plausible that agents do not start their epistemic lives with all the rational epistemic standards they can have.[162] Consider the case of standards that are relevant for religious beliefs. One needs to acquire the concept of religious authority before being able to

[161]See my paper titled "Should Agents Be Immodest?" (m.s. d) for more details.
[162]Elsewhere, me and David Montminy have considered the case of "initial" standards. See Daoust and Montminy (2017; m.s.) for details.

entertain standards such as "trust the religious authorities." Since agents do not necessarily start their epistemic lives with such concepts, the standard "trust the religious authorities" can be acquired later in an agent's epistemic life (e.g., after the agent acquires the relevant concepts).[163]

Now, with respect to acquired epistemic standards, consider the following cases:

Kate and Brad at t_0. At time t_0, Kate thinks that she has no reason to prefer the standard "trust the religious authorities" over the standard "do not trust the religious authorities." Even if she thinks that she has no reason to prefer one standard over the other, she decides to adopt the standard "trust the religious authorities." Brad decides to adopt the standard "do not trust the religious authorities".

Kate and Brad at t_1. After Kate adopts the standard "trust the religious authorities", something happens to her. She suddenly thinks that trusting the religious authorities is more likely to be accurate, even if she hasn't acquired new evidence between t_0 and t_1. She suddenly thinks that, from an accuracy perspective, Brad's standard is suboptimal.

A change in intuitions between t_0 and t_1 can explain why Kate no longer believes that she has no reason to prefer the standard "trust the religious authorities" over the standard "do not trust the religious authorities." However, either (i) such a change in intuitions affects Kate's evidence or (ii) Kate ought to change some of her attitudes without receiving new evidence. Either way, we face a problem. Here is why.

Provided that acquiring epistemic standards changes our intuitions concerning other standards, we can wonder if such a change in intuitions affects an agent's evidence. First, assume that such a change in intuitions affects an agent's evidence. If acquiring epistemic standards changes our intuitions concerning these standards, agents with different epistemic standards *do not share all relevant evidence*. This violates the assumption that agents are epistemic peers. Recall that this chapter is concerned with cases where two agents who *share all relevant evidence* disagree.

In view of the foregoing, Schoenfield probably means that such a change in intuitions does not affect an agent's evidence. But even in making such an assumption, we face a serious

[163]Schoenfield also seems to think that agents acquire epistemic standards over time. If irrelevant influences (such as growing up in one particular community rather than in another) lead agents to adopt distinct incompatible epistemic standards, this means that agents adopt some new epistemic standards over time. In other words, they do not start their epistemic lives with all the epistemic standards they can have.

difficulty. Recall that one motivation in favour of Diachronic Prohibition is that agents should not change their doxastic attitudes without getting new evidence. However, on the assumption that we can appraise epistemic standards differently without acquiring new evidence, Kate ends up changing some of her doxastic attitudes without getting new evidence. Indeed, at time t_0, Kate thinks that she has no reason to prefer one standard over the other. However, at t_1, she is required to believe that her standard is uniquely optimal, and so it would be inconsistent for her to believe that "she has no reason to prefer one standard over the other." In other words, she has to abandon her initial belief that "she has no reason to prefer one standard over the other" at t_1. However, Kate did not acquire new evidence between t_0 and t_1 and her change in perception does not affect her evidence. This means that Kate ends up dropping her belief that she has no reason to prefer one standard over the other without having acquired new evidence. Consequently, assuming that acquiring epistemic standards does not affect an agent's evidence, we are sometimes required to change our doxastic attitudes without getting new evidence. Either way, Schoenfield's argument raises concerns when it comes to acquired epistemic standards.

The third problem concerns the assumption that epistemically rational agents always ought to believe that their standards are more truth-conducive than any other set of standards. Beliefs concerning one's accuracy can be treated like any other beliefs. One can have (or lack) good evidence for or against the truth-conduciveness of one's epistemic standards. Accordingly, if agents have clear evidence that other agents with other standards are equally reliable (or if they lack good evidence that their own standards are more reliable), they should refrain from believing that their epistemic standards are more truth-conducive.

Here is why. To begin with, consider once again Titelbaum and Kopec's Reasoning Room. In fact, suppose that Kate and Brad are in the Reasoning Room. If Schoenfield is right, no matter what kind of information Kate and Brad are provided, they will never believe that they find themselves in such a situation, since agents in the Reasoning Room are equally reliable. If Kate and Brad believe that they find themselves in the Reasoning Room, they believe that distinct incompatible epistemic standards are equally optimal. But this contradicts the claim that agents should believe that their own standards are *more* truth-conducive. So, in

accordance with Strict Immodesty, Brad and Kate will deny that they can find themselves in the Reasoning Room (even if, in fact, they could find themselves in such a situation).

Now, consider the following revised version of the Reasoning Room:

Daily Reasoning Room. Every day, Kate and Brad stand in a room with 8 other people and are given 100 hypotheses to evaluate. Each person in the room possesses the same total evidence relevant to those hypotheses, but each person has distinct incompatible rational epistemic standards. After the participants have evaluated the hypotheses, a great number of independent and extremely reliable brain scanners reveal the following: every participant has formed 90 true beliefs and 10 false beliefs. This result is revealed to the participants day after day. The participant also receive evidence which, relative to their standards, supports the conclusion that the scanners are extremely reliable.[164]

Here, it is patently clear that, day after day, agents have consistent evidence that their standards are not more truth-conducive than others. But if Schoenfield is right, the kind of evidence provided by the reliable brain scanners is not relevant. Following Strict Immodesty, epistemically rational agents should take their standards to be the most truth-conducive ones. So, in the above case, agents should stand their ground and keep believing that their standards are more truth-conducive than others. I find this result implausible: in order to discard the information provided by a great number of independent brain scanners, Kate has to be *overconfident* that her standards are epistemically superior. Being strictly immodest would be irrational given her evidence and standards.

Of course, in the Daily Reasoning Room, agents do not have *independent evidence* for the conclusion that their standards are as reliable as others. Indeed, the scanners provide evidence that agents in the Daily Reasoning Room are equally reliable insofar as agents entertain an epistemic standard such as "trust the brain scanners." However, the issue is not whether the scanners provide independent evidence for the conclusion that agents are equally reliable. The issue is whether an agent's rational epistemic standards will recommend not trusting the information provided by the scanners. Note, for instance, that the participant receive evidence which, relative to their standards, supports the conclusion that the scanners are extremely reliable. In the Daily Reasoning Room, not trusting the information provided by

[164]Note that the scanners could reveal other relevant information to participants. For instance, some participants could falsely believe that they have failed to satisfy their own epistemic standards. The scanners could reveal to them that they have satisfied their own epistemic standards.

the scanners amounts to being overconfident. Accordingly, agents with rational epistemic standards will trust the scanners.

This leads me to conclude that Kate's beliefs concerning the truth-conduciveness of her standards can be confirmed or disproved by her evidence. If she lacks sufficient evidence to believe that her standards are more truth-conducive, she should not believe it. So, it is false that epistemically rational agents ought to believe that their standards are more truth-conducive than others. An epistemically rational agent ought to believe what his or her evidence supports, and the evidence might not support the belief that that his or her standards are epistemically superior.

5.4.3. The Objection from Dutch Books

Another common explanation of why there seems to be diachronic norms prohibiting an agent from changing his or her epistemic standards is that such an agent would display a type of diachronic incoherence, and would then be subject to a diachronic Dutch Book. Roughly, Dutch Book arguments aim at establishing epistemic norms from the fact that, if some epistemic behaviours are permitted, then an agent could be subject to a sure-loss combination of bets. These arguments presuppose that an agent's degrees of belief in some propositions are in line with his or her dispositions to accept some bets concerning these propositions.[165]

Consider the case of prior functions in Bayesian epistemology. Imagine that Kate currently believes that, on the condition that she obtains evidence E, she would assign a credence of X in P. However, suppose that she can also refuse to endorse certain of her possible future judgments. For instance, once Kate will acquire evidence E, she could assign a credence of Y in P such that $X \neq Y$. This can be explained by the fact that Kate decided to *change her prior function*. However, if Kate is permitted to change her prior function over time, she would be vulnerable to a diachronic combination of bets which, taken collectively, will lead her to a sure loss. Therefore, it seems that there are diachronic norms prohibiting

[165]See Vineberg (2016) for an overview of the debates surrounding the Dutch Book methodology.

one's from arbitrarily changing one's epistemic standards over time, since agents who arbitrarily change their epistemic standards over time are vulnerable to a Dutch Book.

The Dutch Book methodology has clear limits: the assumption that an agent's credences in various propositions line up with his or her dispositions to accept some bets concerning these propositions is contentious. Also, it is far from clear that diachronic incoherence is necessarily irrational. For instance, there are situations where it becomes patently clear that prior probabilities are mistaken. In such cases, revising prior probabilities is the right thing to do, even if agents display diachronic incoherence or are vulnerable to a Dutch Book. In the context of statistical modelling, Cox argues that:

> [temporal incoherency] is not inevitable and there is nothing intrinsically inconsistent in changing prior assessments, in particular in the light of experience obtained either in the process of data collection or from the data themselves.... The play of chance may have been unkind. The data may be contaminated. The prior may be based on a misconception (Cox 2006, 78).

Nevertheless, for the sake of the argument, let's assume for a moment that the Dutch Book methodology is conclusive and that diachronic incoherence is epistemically irrational. Provided that there is something wrong with arbitrarily changing one's attitudes or prior function (as the Dutch Book methodology suggests), this does not imply that when an epistemically rational agent adopts an attitude at time t_0, there is a diachronic norm prohibiting him or her from adopting incompatible attitudes at time t_1 (as Diachronic Prohibition suggests). In other words, Diachronic Prohibition is *compatible* with the Dutch Book methodology, but it is not *supported* by the Dutch Book methodology.

Here is why we can assume that there is something wrong with arbitrarily changing one's prior function and, at the same time, that there is no diachronic norm prohibiting one from changing his prior function over time. There could be something wrong with arbitrarily changing one's prior function because there is a *unique* rational prior function. If there is a unique rational prior function, then epistemically rational agents cannot arbitrarily change their prior function: if they change their rational prior function, they would necessarily end up with a less than fully rational prior function. Hence, provided that there is a unique rational

prior function, there is no need to explain what is wrong with arbitrarily changing one's priors by appealing to Diachronic Prohibition.

To be clear: my point here is that, even if we accept the Dutch Book methodology, there are different explanations of why arbitrarily changing one's prior function is problematic. Since some of them do not support Diachronic Prohibition, we can't conclude that Diachronic Prohibition is true.

Suppose the Dutch Book methodology is correct. Then, there are at least two explanations of why there is something wrong with arbitrarily changing one's prior function. First, perhaps there is a unique rational prior function and Diachronic Prohibition is false. Second, perhaps there is no unique rational prior function, but Diachronic Prohibition is true. As we can see, one of these explanations denies Diachronic Prohibition. So even if we accept the Dutch Book methodology, this does not support Diachronic Prohibition: the possibility of diachronic Dutch Books is compatible with the denial of Diachronic Prohibition.[166]

This should come as no surprise to those who are familiar with the exchange between White and Meacham. White (2005, 454–55) has argued that, since there is something wrong with arbitrarily changing one's prior function, permissiveness must be false. According to White, since agents should not arbitrarily change their prior function, there must be a unique rational prior function. Meacham (2014, 1206) roughly replied that Diachronic Prohibition can explain why there is something wrong with arbitrarily changing one's prior function.[167] According to Meacham, since Diachronic Prohibition is compatible with the Permissive Epistemic Standards Thesis, thinking that there is something wrong with arbitrarily changing one's prior function does not commit us to the denial of permissiveness. However, note that Meacham does not deny White's explanation: he merely highlights that there is an alternative explanation (Diachronic Prohibition) which is compatible with permissiveness. What is clear

[166]Objection: even if there is a uniquely rational prior distribution, we can imagine that an agent's prior does not line up with the uniquely rational prior function. Following the Dutch Book methodology, he or she should not change his or her credence distribution. So, Diachronic Prohibition is supported by the Dutch Book Methodology. Response: under the assumption that there is a uniquely rational prior function, an agent who doesn't entertain such a function is not ideally rational, and this is not the kind of agent I am concerned with in this paper.

[167]Specifically, Meacham argued that Bayesian conditionalization, which implies Diachronic Prohibition, can explain why epistemically rational agents should not arbitrarily change their prior function.

in the above exchange is that, even if the Dutch Book methodology shows that there is something wrong with arbitrarily changing one's prior function, there are different ways to explain such a result. Some of these explanations support Diachronic Prohibition while others do not. Hence, the Dutch Book methodology neither supports nor disproves Diachronic Prohibition.[168]

5.5. Conclusion: Back to the Normativity of Epistemic Rationality.

In this chapter, I argued against the Permissive Epistemic Standards Thesis (at least with respect to ideal theories of epistemic rationality). Using Condorcet's Jury Theorem, I argued that putative equally reliable but incompatible epistemic standards entail the existence of a unique and more reliable epistemic standard.

The argument developed in this chapter was the last step towards vindicating the normativity of epistemic rationality. Here is why. In chapter 1, I have argued that the Minimal Normative Hypothesis and the Modest Reductionist Hypothesis are sufficient for vindicating the normativity of epistemic rationality. In chapter 2, I have derived the deontic significance of apparent reasons from the deontic significance of reasons. Indeed, my starting point was that epistemic reasons are deontically significant, which led me to the conclusion that apparent reasons are also deontically significant. This allowed me to confirm the Minimal Normative Hypothesis:

Minimal Normative Hypothesis. Agents ought to respond correctly to (apparent) sufficient epistemic reasons they have. Responding correctly to (apparent) reasons one has is deontically significant.

In other words, substantive requirements of epistemic rationality are normative. Where does that leave us with respect to the normativity of epistemic rationality *in general*? We know that, if reductionism is true, then only substantive requirements play a distinct explanatory role in the theory of rationality. We know that those requirements are normative. So, we now simply need to determine whether reductionism is true.

[168]The connection between prohibiting diachronic incoherence and permissiveness has also been discussed by Hedden (2015b, 717–18) and Titelbaum and Kopec (m.s., 9).

In the conclusion of chapter 3, I have argued that, when compared with substantive requirements of epistemic rationality, putative structural requirements such as Inter-Level Coherence or Intra-Level Coherence have no distinct explanatory role in the theory of epistemic rationality. The last structural requirement that could play an explanatory role in the theory of epistemic rationality would be Consistency. This brought us to chapter 4, where I have argued that Consistency plays an explanatory role in the theory of epistemic rationality only to the extent that there are situations in which rational agents can entertain distinct incompatible epistemic standards. However, in the present chapter, I have argued that, when it comes to ideally rational agents, the Superior Standard Thesis is true. According to such a view, putative equally reliable standards always collapse into a unique, more reliable standard. This means that the Permissive Epistemic Standards Thesis is false for ideally rational agents.

Thus, Consistency plays no explanatory role in ideal theories of epistemic rationality. When combined, the arguments put forth in chapters 3 to 5 support the Modest Reductionist Hypothesis, namely:

Modest Reductionist Hypothesis. In ideal theories of epistemic rationality, putative structural requirements such as Inter-Level Coherence, Intra-Level Coherence or Consistency have no distinct explanatory role when compared with substantive requirements of epistemic rationality.

Thus, in ideal theories of epistemic rationality, requirements of reasons-responsiveness are deontically significant and structural requirements play no distinct explanatory role. We have reached our destination. Ideal epistemic rationality is deontically significant.

What about non-ideal theories of epistemic rationality? Perhaps reductionism fails with respect to such theories, in the sense that Consistency would play a distinct explanatory role in such theories. However, this doesn't mean that non-ideal theories of epistemic rationality are not normative. Perhaps Consistency is normative. I have briefly explored this possibility at the beginning of chapter 4. So, with respect to non-ideal theories of epistemic rationality we are left with three possibilities, namely:

(1) Consistency plays no distinct explanatory role in non-ideal theories of epistemic rationality, and those theories are normative;

(2) Consistency plays a distinct explanatory role in non-ideal theories of epistemic rationality, but those theories are still normative e.g., there is an explanation of why structural requirements like Consistency are normative);

(3) Non-ideal theories of epistemic rationality are not normative.

Determining which possibility is correct is a task for another day.

Conclusion (English Version)

Summary of the Thesis

I have argued that epistemic rationality is normative, or more specifically, that agents ought to be epistemically rational. My thesis relies on four key assumptions. First, there are some properly epistemic goods such as truth. Second, there are substantive requirements of epistemic rationality, such as responding correctly to (apparent) epistemic reasons one has. Third, the weight of epistemic reasons can be represented with epistemic probabilities. Fourth, in the right conditions, if A is ideally rational, A satisfies available epistemic standards that optimize his or her ratio of true to false beliefs.

There are two steps for my argument in favour of the normativity of epistemic rationality. First, in chapter 2, I have argued that the normativity of apparent reasons is tied to the normativity of reasons. Since the normativity of reasons is fairly uncontroversial, it should also be uncontroversial that apparent reasons are normative. Following the second assumption of my thesis, there are substantive requirements of epistemic rationality, such as responding correctly to (apparent) epistemic reasons one has. In view of the foregoing, those requirements bear deontic significance.

Now, in addition to substantive requirements of epistemic rationality, there are putative structural requirements of epistemic rationality (such as Consistency, Intra-Level Coherence or Inter-Level Coherence). It is unclear whether those requirements are normative. Besides, most of Kolodny's arguments against the normativity of epistemic rationality have to do with structural requirements. Therefore, I had to explain why these requirements are no obstacle to the normativity of epistemic rationality. My strategy has been to "eliminate" these requirements, or to argue that they play no distinct explanatory role in ideal theories of epistemic rationality.

In chapter 3, I have argued that two structural requirements (Inter-Level Coherence and Intra-Level Coherence) play no distinct explanatory role in the theory of epistemic rationality. Roughly, I have offered the following argument:

(1) The Closure Principle roughly states that, if A is rational to believe that P is A is rational to believe that Q, then A is rational to believe that (P^Q).

(2) The Closure Principle is either true or false.

(3) If the Closure Principle is true, there is an explanation (in terms of reasons) of why agents ought to satisfy structural requirements of epistemic rationality—agents who violate structural requirements fail to respond correctly to their reasons (Reisner 2011).

(4) If the Closure Principle is false, Inter-Level Coherence and Intra-Level Coherence are not genuine requirements of epistemic rationality.

(C) Therefore, whether the Closure Principle is true or false, Inter-Level Coherence and Intra-Level Coherence play no distinct explanatory role in the theory of epistemic rationality.

Chapter 3 remained neutral on whether Consistency plays a distinct explanatory role in the theory of epistemic rationality. In chapter 4, I have argued that Consistency plays a distinct explanatory role in epistemically permissive situations. In such situations, a rational agent is epistemically permitted to believe that P and also to believe that ~P (relative to a body of epistemic reasons). Permissiveness is a live possibility when it comes to epistemic standards. For example, agents might be rationally permitted to entertain standards X, but also to entertain standards Y (where standards X and standards Y are incompatible with each other).

However, in chapter 5, I have argued that, when it comes to the ideal theories of rationality, this type of permissiveness is false. My argument relies on Condorcet's Jury Theorem. Assuming that there are distinct incompatible but equally reliable epistemic standards St_1, St_2, ..., St_n, there exists a unique meta-standard that is more reliable than St_1, St_2, ..., St_n. Following the fourth assumption of my thesis, if A is ideally rational, A satisfies available epistemic standards that optimize his or her ratio of true to false beliefs. So, ideally rational agents with the same epistemic reasons satisfy uniquely determined epistemic

standards, and permissiveness concerning epistemic standards is false for ideally rational agents.

The arguments of chapters 3 to 5 entail that structural requirements of rationality play no distinct explanatory role in ideal theories of epistemic rationality. This allowed me to offer the following argument in favour of the normativity of ideal epistemic rationality:

(1) There are two types of requirements of epistemic rationality, namely, structural requirements and substantive requirements.

(2) Substantial requirements of epistemic rationality are normative.

(3) Structural requirements play no distinct explanatory role in ideal theories of epistemic rationality.

(C) So, ideal theories of epistemic rationality are normative.

Novel Aspects of the Thesis

I have offered new arguments and perspectives on the nature and the normativity of epistemic rationality. In chapter 1, I have rejected Broome's argument in favour of structuralism about epistemic rationality. His arguments are inconclusive and do not allow us to make sense of ordinary attributions of irrationality. If rationality merely consists in being coherent, dogmatic agents, skeptics, conspiracy theorists and grue-projectors can be fully rational insofar as they remain coherent. This conflicts with ordinary attributions of epistemic irrationality.

My theory of epistemic rationality also differs from the ones developed by Errol Lord and Benjamin Kiesewetter. First, Lord and Kiesewetter do not respond to some of the objections put forth by Niko Kolodny. According to Kolodny, there should be a clear connection between epistemic norms and final epistemic value (e.g., truth). For example, X is a bad epistemic norm if one can satisfy X while only having false beliefs. As I indicated in section 1.2.5, this objection undermines the normativity of apparent reasons (or of reasons that consist in facts about appearances). Kiesewetter's and Lord's vindications of the normativity

of epistemic rationality do not provide a satisfactory response to Kolodny's objection. Chapter 2 offers a new vindication for the normativity of apparent reasons (or of reasons that consist in facts about appearances) and responds to Kolodny's objection.

Also, Kiesewetter and Lord think that structural requirements play no distinct explanatory role in the theory of epistemic rationality. I agree with this verdict (at least when it comes to ideal theories of epistemic rationality). However, their accounts of epistemic rationality presuppose uniqueness. As I explained in chapter 4, this assumption is ill-founded. So, they do not provide a satisfactory argument for reductionism. In fact their reductionism about structural requirements of epistemic rationality is question-begging. This is why, in chapters 4 and 5, I shed light on permissiveness and its implications.

Finally, my own conclusions are weaker than Kiesewetter's and Lord's. The reductionism I endorse does not necessarily apply to the non-ideal theories of epistemic rationality. As I explained in chapters 2 and 5, there is no direct relationship between ideal theories and non-ideal theories of epistemic rationality. For instance, there could be second-best problems which explain why non-ideal theories should not merely approximate ideal theories. Thus, as I explained at the end of chapter 5, robust reductionism (which covers ideal and non-ideal theories of epistemic rationality) might be false.

Limits of the Argument and Future Work

My argument does not generalize easily to non-ideal theories of epistemic rationality. Ideal and non-ideal theories of epistemic rationality have different features. A good argument in favour of the normativity of epistemic rationality must take these particularities into account. That having being said, this "generalization" problem is not specific to theories of epistemic rationality. Rather, it is a consequence of the epistemic theory of the second best, and any idealized theory of epistemic norms might not generalize well to its non-ideal counterpart.

Also, as I indicated earlier, my argument rests on four assumptions, namely the existence of properly epistemic goods like truth, the existence of substantive requirements of

epistemic rationality, the probabilistic representation of the weight of epistemic reasons and the reliability criterion. Any philosophical argument relies on some basic assumptions. Nevertheless, it would be appropriate to consider different starting points for vindicating the normativity of epistemic rationality. For example, can we argue that epistemic rationality is normative without assuming any connection between reliability and rationality, or without assuming that the weight of epistemic reasons can be represented with epistemic probabilities?

Furthermore, I have focused on three structural requirements of epistemic rationality, namely, Consistency, Inter-Level Coherence and Intra-Level Coherence. Other putative structural requirements can be found in the contemporary literature on epistemic rationality. For example, some putative requirements of epistemic rationality govern combinations of credences rather than combinations of beliefs. Thus, the reductionism discussed in chapters 3 to 5 leaves aside some plausible structural requirements governing credences. Future research will be needed to determine whether reductionism concerning structural requirements also applies to requirements governing credences.

Finally, I did not address the normativity of practical rationality. The requirements of practical rationality differ from the requirements of epistemic rationality in many ways. Here is a simple example: permissivism. I have argued that, when it comes to ideal theories of epistemic rationality, extreme permissiveness is false. However, when it comes to theories of practical rationality (including ideal theories of practical rationality), permissivism is true. Suppose that Buridan's ass is planning what to intend when facing identical stacks of hay. Call these options Left and Right. Suppose it has all the relevant information to make a good decision. It would be irrational for Buridan's ass to do nothing and starve to death. Yet, it lacks sufficient reason to choose Left rather than choosing Right (and *vice versa*). So, the donkey is in a permissive situation. Both options are rationally permitted but incompatible with each other.

Extreme permissivism is a notable difference between epistemic rationality and practical rationality. Accordingly, the arguments developed in chapters 2 to 5 do not apply

directly to practical rationality. Thus, future research will be needed to determine whether structural requirements of practical rationality are normative.

Conclusion

Résumé de la thèse

J'ai soutenu la thèse selon laquelle la rationalité épistémique est normative, ou plus précisément, que les agents doivent être épistémiquement rationnels. Ma thèse repose sur quatre présupposés importants que je n'ai pas démontrés (et que je ne sais pas comment démontrer). Un : il existe des « biens épistémiques » comme la vérité. Deux : la rationalité comprend des exigences substantielles, comme le fait de répondre correctement à ses raisons raisons (apparentes). Trois : le poids ou la force des raisons peut être représenté par des probabilités épistémiques. Quatre : dans les bonnes conditions, si A est idéalement rationnel, A satisfait des standards épistémiques disponibles qui optimisent son ratio de croyances vraies à fausses.

Après avoir exposé ces présupposés, j'ai procédé en deux étapes pour démontrer la normativité de la rationalité. D'abord, dans le chapitre 2, j'ai montré que la normativité des raisons apparentes est inséparable de la normativité des raisons. Puisqu'il y a un vaste consensus entourant la normativité des raisons, il est tout aussi plausible que les raisons apparentes soient normatives. Suivant le second présupposé de ma thèse, cela signifie qu'une partie intégrante de la rationalité épistémique est normative, soit les exigences substantielles stipulant que les agents doivent répondre correctement à leurs raisons épistémiques (apparentes).

Cela dit, il existe des exigences structurales de la rationalité (notamment, les exigences de cohérence ou d'enkrasie). La normativité de ces exigences est incertaine. D'ailleurs, la plupart des critiques de Kolodny à l'égard de la normativité de la rationalité sont dirigés contre les exigences structurales. Il fallait donc expliquer pourquoi ces exigences ne sont pas un obstacle à la normativité de la rationalité. Ma stratégie a été d'éliminer ces exigences : j'ai défendu l'idée selon laquelle ces exigences ne jouent aucun rôle explicatif distinct dans la théories de la rationalité épistémique idéale.

Dans le chapitre 3, j'ai défendu ce réductionnisme pour deux exigences structurales (*Inter-Level Coherence* et *Intra-Level Coherence*). En gros, j'ai offert le raisonnement suivant :

(1) Selon le principe de clôture, si A est rationnel de croire que P est A est rationnel de croire que Q, alors A est rationnel de croire que (P^Q).

(2) Le principe de clôture est vrai ou faux.

(3) Si le principe de clôture est vrai, il existe une explication (en termes de raisons) de pourquoi les agents doivent satisfaire les exigences structurales de la rationalité. En effet, les agents qui enfreignent les exigences structurales de la rationalité répondent incorrectement à leurs raisons (Reisner 2011).

(4) Si le principe de clôture est faux, *Inter-Level Coherence* et *Intra-Level Coherence* ne sont pas des exigences de la rationalité épistémique.

(C) Donc, que le principe de clôture soit vrai ou faux, *Inter-Level Coherence* et *Intra-Level Coherence* ne jouent pas de rôle explicatif distinct dans la théorie de la rationalité.

Le chapitre 3 est demeuré neutre quant au réductionnisme pour l'exigence de non-contradiction (*Consistency*). Selon cette exigence, un agent rationnel doit éviter de croire que P et de croire que ~P simultanément. J'ai défendu l'idée selon laquelle l'exigence de non-contradiction joue un rôle explicatif distinct dans les situations épistémiquement permissives. Dans de telles situations, un agent rationnel peut croire que P et peut aussi croire que ~P (relativement à un ensemble de raisons épistémiques). Naturellement, il serait surprenant que, relativement à un ensemble de raisons épistémiques *et de standards épistémiques*, un agent ait de telles permissions. Mais cela ne signifie pas que la rationalité épistémique est impermissive. En effet, il se pourrait que différents agents rationnels puissent satisfaire différents standards épistémiques. Par exemple, il se pourrait qu'un agent puisse être rationnellement permis d'avoir les standards X aussi bien que les standards Y (bien que les standards X et Y soient incompatibles entre eux).

Dans le chapitre 5, j'ai cependant soutenu que, lorsqu'il est question des théories idéales de la rationalité, ce type de permissivisme est forcément faux. Ma stratégie dans ce chapitre repose largement sur le théorème du jury de Condorcet. À supposer qu'il existe des standards épistémiques incompatibles $\{St_1, St_2, ..., St_n\}$ mais également fiables, il existe forcément un « méta-standard » unique et plus fiable que $\{St_1, St_2, ..., St_n\}$. Étant donné le quatrième présupposé de ma thèse (les agents idéalement rationnels satisfont les standards qui optimisent leur ratio de croyances vraies à fausses), les agents idéalement rationnels satisfont tous le même ensemble de standards épistémiques. En d'autres termes, le permissivisme des standards épistémiques est faux pour les agents idéalement rationnels.

Ainsi, les arguments des chapitres 3 à 5 nous permettent alors de conclure que les exigences structurales de rationalité ne jouent aucun rôle explicatif distinct dans les théories de la rationalité épistémique idéale. Cela me permet donc d'offrir l'argument suivant en faveur de la normativité de la rationalité épistémique idéale :

(1) Il existe deux familles d'exigences de la rationalité épistémique, soit les exigences structurales et les exigences substantielles.

(2) Les exigences substantielles de la rationalité épistémique sont normatives.

(3) Les exigences structurales ne jouent aucune ne jouent aucun rôle explicatif distinct dans les théories idéales de la rationalité épistémique.

(C) Donc, les théories de la rationalité épistémique idéale sont normatives.

Aspects distinctifs de la thèse

Ma thèse se démarque de toutes les contributions existantes sur le sujet. Dans le chapitre 1, j'ai rejeté les arguments en faveur du structuralisme offerts par John Broome. Non seulement ces arguments sont non-concluants, mais ils ne nous permettent pas de rendre compte des attributions ordinaires d'irrationalité. Si la rationalité consistait simplement dans le fait d'être cohérent, n'importe quel individu dogmatique, dupe, sceptique ou conspirationniste

pourrait être épistémiquement rationnel *en autant qu'il soit cohérent*. Cette approche est en contradiction directe avec les attributions ordinaires d'irrationalité épistémique.

Ma thèse me démarque aussi des approches basées sur les exigences substantielles développées par Errol Lord et Benjamin Kiesewetter. D'une part, leurs théories ne prennent pas au sérieux certaines objections développées par Niko Kolodny. Un des arguments de Kolodny est qu'il devrait y avoir une connexion claire entre les normes épistémiques et la valeur épistémique finale (e.g., la vérité). Par exemple, X n'est pas une bonne norme épistémique si l'on peut satisfaire X tout en ayant seulement des croyances fausses. Comme je l'ai indiqué au point 1.2.5, cette objection remet sérieusement en cause la normativité des raisons apparentes (ou des raisons comprises au sens large, incluant les faits à propos des apparences). Le fait d'inclure les faits concernant les apparences dans les raisons d'un agent ne répond pas adéquatement à l'objection de Kolodny. Ma thèse, en particulier le chapitre 2, offre une nouvelle justification de la normativité des raisons apparentes et répond directement à l'objection de Kolodny.

D'autre part, Kiesewetter et de Lord soutiennent que les exigences structurales ne jouent aucun rôle explicatif distinct dans la théorie de la rationalité épistémique. Je suis d'accord avec leur conclusion (à tout le moins, en ce qui concerne les théories idéales de la rationalité épistémiques). Or, leurs théories présupposent une forme de non-permissivisme épistémique. Comme je l'ai expliqué au chapitre 4, ce présupposé est insatisfaisant, surtout si l'on souhaite offrir une défense satisfaisante du réductionnisme quant aux exigences structurales. En d'autres termes, leur réductionnisme est une pétition de principe. C'est pourquoi j'ai proposé une analyse du permissivisme et de ses implications dans les chapitres 4 et 5.

Finalement, mes conclusions sont beaucoup plus modestes que celles de Kiesewetter et Lord. Le réductionnisme que je défends ne s'applique pas forcément aux théories non-idéales de la rationalité. Comme je l'ai expliqué dans les chapitres 2 et 5, il n'y a pas forcément de relation directe entre les théories idéales et les théories non-idéales de la rationalité. Cela peut être dû à un problème d'optimum de second rang (*second-best problem*). Comme je l'ai

expliqué dans la seconde moitié du chapitre 5, un réductionnisme complet quant aux exigences structurales me semble injustifié, puisque le permissivisme pourrait être vrai dans les théories non-idéales de la rationalité épistémique.

Limites de la thèse et travaux futurs

Comme je viens de l'indiquer, cette thèse a le désavantage de ne pas être généralisable aux théories non-idéales de la rationalité épistémique. Les théories idéales de la rationalité épistémique ont leurs particularités, tout comme les théories non-idéales. Un bon argument en faveur de la normativité de la rationalité épistémique doit prendre ces particularités en compte. Cela dit, ce n'est pas un défaut propre à la théorie de la rationalité épistémique. Il s'agit plutôt une conséquence de la théorie épistémique de l'optimum de second rang. En d'autres termes, toute théorie épistémologique idéalisée pourrait n'avoir aucune répercussion sur sa contrepartie non-idéale.

Aussi, comme je l'ai indiqué il y a quelques paragraphes, cette thèse repose sur quatre principes importants, soit l'existence de « biens épistémiques » comme la vérité, le non-structuralisme concernant les exigences de la rationalité épistémique, la représentation probabiliste du poids des raisons épistémiques et le critère de fiabilité. Naturellement, toute recherche admet des présupposés. Ils sont le point de départ nécessaire de toute enquête. Néanmoins, il serait opportun d'étudier la normativité de la rationalité épistémique à partir de points de départ différents dans des recherches futures. Par exemple, peut-on montrer que la rationalité épistémique est normative tout en abandonnant le critère de fiabilité, ou en optant pour une représentation différente du poids des raisons?

Troisièmement, je n'ai pas analysé toutes les exigences structurales que l'on retrouve dans la littérature contemporaine sur la rationalité épistémique. Par exemple, je me suis concentré sur les exigences gouvernant les croyances, et non sur les exigences qui gouvernent les degrés de confiance (*credences*). Ainsi, le réductionnisme discuté dans les chapitres 3 à 5 a fait l'économie de certaines exigences structurales plausibles gouvernant les degrés de

confiance. Des recherches futures seront nécessaires pour déterminer si la thèse du réductionnisme s'applique aussi à ces exigences.

Finalement, cette thèse est demeurée neutre quant à la normativité de la rationalité pratique. Les normes de rationalité pratique diffèrent des normes de rationalité épistémique. Prenons un exemple simple : la question du permissivisme. Dans cette thèse, j'ai soutenu que le permissivisme extrême était faux dans les théories idéales de la rationalité épistémique. Or, le permissivisme est assurément vrai dans les théories de la rationalité pratique (et ce, incluant les théories idéales de la rationalité pratique). Pensons, par exemple, à l'âne de Buridan. Imaginons qu'un âne affamé soit à égale distance de deux bottes de foin identiques à toutes fins pratiques (le foin des deux bottes a le même goût, les deux bottes sont de la même taille, et ainsi de suite). L'âne délibère quant à savoir s'il devrait manger la botte de foin à sa gauche ou celle à sa droite. Imaginons de plus que l'âne possède toutes les informations pertinentes pour faire un choix éclairé. Dans un tel contexte, il n'y a pas de raison concluante pour l'âne d'aller vers une botte de foin plutôt que vers l'autre. Ainsi, l'âne se trouve dans une situation permissive : il est rationnel pour l'âne de manger la botte de foin à sa gauche, et il est aussi rationnel pour l'âne de manger la botte de foin à sa droite, même si ses options sont incompatibles entre elles.

Le permissivisme extrême est une différence parmi d'autres entre la rationalité épistémique et la rationalité pratique. C'est pourquoi les raisonnements et arguments développés dans les chapitres 2 à 5 ne peuvent pas s'appliquer directement au cas de la rationalité pratique. Ainsi, des recherches futures seront nécessaires afin de déterminer si la rationalité pratique est normative.

Bibliographie

Ahlstrom-Vij, Kristoffer, and Jeffrey Dunn. 2014. “A Defence of Epistemic Consequentialism.” *The Philosophical Quarterly* 64 (257): 541–51. https://doi.org/10.1093/pq/pqu034.

Alexander, David J. 2013. “The Problem of Respecting Higher-Order Doubt.” *Philosopher’s Imprint* 13 (18): 1–12.

Allais, Maurice. 1953. “Le comportement de l’homme rationnel devant le risque: Critique des Postulats et axiomes de l’école américaine.” *Econometrica* 21 (4): 503–46.

Atkinson, David, and Jeanne Peijnenburg. 2006. “Probability without Certainty: Foundationalism and the Lewis–Reichenbach Debate.” *Studies in History and Philosophy of Science Part A* 37 (3): 442–453.

Atkinson, David, and Jeanne Peijnenburg. 2009. “Justification by an Infinity of Conditional Probabilities.” *Notre Dame Journal of Formal Logic* 50 (2): 183–93. https://doi.org/10.1215/00294527-2009-005.

Bachrach, Yoram, Thore Graepel, Gjergji Kasneci, Michal Kosinski, and Jurgen Van Gael. 2012. “Crowd IQ: Aggregating Opinions to Boost Performance.” In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 535–542.

Ballantyne, Nathan, and E. J. Coffman. 2011. “Uniqueness, Evidence, and Rationality.” *Philosopher’s Imprint* 11 (18): 1–13.

Bélanger, Marco. 2011. *Existe-t-il des dilemmes moraux insolubles?* Paris: Éditions L’Harmattan.

Berker, Selim. 2013a. “Epistemic Teleology and the Separateness of Propositions.” *Philosophical Review* 122 (3): 337–393.

Berker, Selim. 2013b. “The Rejection of Epistemic Consequentialism.” *Philosophical Issues* 23 (1): 363–387.

Berker, Selim. 2015. “Reply to Goldman: Cutting up the One to Save the Five in Epistemology.” *Episteme* 12 (02): 145–153.

Bermúdez, José Luis. 2009. *Decision Theory and Rationality*. Oxford: Oxford University Press.

Blanshard, Brand. 1974. *Reason and Belief*. London: Allen and Unwin.

Brady, Michael. 2009. “Curiosity and the Value of Truth.” In *Epistemic Value*, edited by Adrian Haddock, Allan Millar, and Duncan Pritchard, 265–283. Oxford, New York: Oxford University Press.

Briggs, R. 2017. “Normative Theories of Rational Choice: Expected Utility.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2017.

Broome, John. 2005. “Does Rationality Give Us Reasons?” *Philosophical Issues* 15 (1): 321–337.

Broome, John. 2007a. “Does Rationality Consist in Responding Correctly to Reasons?” *Journal of Moral Philosophy* 4 (3): 349–374.

Broome, John. 2007b. “Wide or Narrow Scope?” *Mind* 116 (462): 359–370.

Broome, John. 2008. “Reply to Southwood, Kearns and Star, and Cullity.” *Ethics* 119 (1): 96–108.

Broome, John. 2013. *Rationality through Reasoning*. Oxford: John Wiley & Sons.

Brueckner, Anthony, and Alex Bundy. 2012. “On ‘Epistemic Permissiveness.’” *Synthese* 188 (2): 165–177.

Brunero, John. 2010. “The Scope of Rational Requirements.” *The Philosophical Quarterly* 60 (238): 28–49.

Buchak, Lara. 2013. *Risk and Rationality*. New York and Oxford: Oxford University Press.

Buchak, Lara, and Philip Pettit. 2015. “The Case of Group Agents.” Edited by Iwao Hirose and Andrew Reisner. *Weighing and Reasoning: Themes from the Philosophy of John Broome*, 207–31.

Bykvist, Krister, and Anandi Hattiangadi. 2007. “Does Thought Imply Ought?” *Analysis* 67 (296): 277–85. https://doi.org/10.1111/j.1467-8284.2007.00691.x.

Buchak, Lara. 2013. *Risk and Rationality*. New York and Oxford: Oxford University Press.

Buchak, Lara. 2014. “Belief, Credence, and Norms.” *Philosophical Studies* 169 (2): 285–311.

Caie, Michael. 2013. “Rational Probabilistic Incoherence.” *Philosophical Review* 122 (4): 527–575.

Carr, Jennifer Rose. 2015. “Don’t Stop Believing.” *Canadian Journal of Philosophy* 45 (5–6): 744–766.

Chang, Ruth, ed. 1997. *Incommensurability, Incomparability, and Practical Reason*. Cambridge (MA): Harvard University Press.

Chang, Ruth. 2001. “Against Constitutive Incommensurability or Buying and Selling Friends.” *Philosophical Issues* 11 (1): 33–60.

Cherniak, Christopher. 1990. *Minimal Rationality*. Cambridge (MA): Mit Press.

Chislenko, Eugene. 2014. “Moore’s Paradox and Akratic Belief.” *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12127.

Christensen, David. 1992. “Confirmational Holism and Bayesian Epistemology.” *Philosophy of Science* 59 (4): 540–557.

Christensen, David. 1994. “Conservatism in Epistemology.” *Noûs* 28 (1): 69–89.

Christensen, David. 2000. “Diachronic Coherence versus Epistemic Impartiality.” *The Philosophical Review* 109 (3): 349–71. https://doi.org/10.2307/2693694.

Christensen, David. 2004. *Putting Logic in Its Place*. Oxford: Oxford University Press.

Christensen, David. 2009. “Disagreement as Evidence: The Epistemology of Controversy.” *Philosophy Compass* 4 (5): 756–767.

Christensen, David. 2010. “Higher-Order Evidence.” *Philosophy and Phenomenological Research* 81 (1): 185–215.

Christensen, David. 2013. “Epistemic Modesty Defended.” In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey. Oxford: Oxford University Press.

Christensen, David. 2014. “Conciliation, Uniqueness and Rational Toxicity.” *Noûs*. https://doi.org/10.1111/nous.12077.

Clarke, Roger. 2013. “Belief Is Credence One (in Context).” *Philosopher’s Imprint* 13 (11): 1–18.

Cloos, Christopher Michael. 2015. “Responsibilist Evidentialism.” *Philosophical Studies* 172 (11): 2999–3016.

Coates, Allen. 2012. “Rational Epistemic Akrasia.” *American Philosophical Quarterly* 48 (2): 113–24.

Coliva, Annalisa. 2015. *Extended Rationality: A Hinge Epistemology*. New York: Palgrave Macmillan.

Condorcet, Nicolas de. 1976. “Essay on the Application of Mathematics to the Theory of Decision-Making.” In *Condorcet: Selected Writings*. Indianapolis: Bobbs-Merrill.

Conee, Earl. 2010. “Rational Disagreement Defended.” In *Disagreement*, edited by Ted Warfield and Richard Feldman, 69–90. Oxford: Oxford University Press.

Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford University Press.

Côté-Bouchard, Charles. 2015. “Epistemic Instrumentalism and the Too Few Reasons Objection.” *International Journal of Philosophical Studies* 23 (3): 337–55. https://doi.org/10.1080/09672559.2015.1042007.

Côté-Bouchard, Charles. 2016. “Can the Aim of Belief Ground Epistemic Normativity?” *Philosophical Studies* 173 (12): 3181–98. https://doi.org/10.1007/s11098-016-0657-8.

Côté-Bouchard, Charles. 2017. “Is Epistemic Normativity Value-Based?” *Dialogue* 56 (3): 407–30. https://doi.org/10.1017/S0012217317000506.

Cox, David Roxbee. 2006. *Principles of Statistical Inference*. Cambridge: Cambridge University Press.

Daoust, Marc-Kevin. 2017. “Epistemic Uniqueness and the Practical Relevance of Epistemic Practices.” *Philosophia* 45 (4): 1721–33. https://doi.org/10.1007/s11406-017-9867-3.

Daoust, Marc-Kevin. 2018a. “Epistemic Akrasia and Epistemic Reasons.” *Episteme*, 1–21. https://doi.org/10.1017/epi.2018.6.

Daoust, Marc-Kevin. 2018b. “The Explanatory Role of Consistency Requirements” *Synthese*. https://doi.org/10.1007/s11229-018-01942-8.

Daoust, Marc-Kevin. 2018c. “Peer Disagreement and the Bridge Principle.” *Topoi*. https://doi.org/10.1007/s11245-018-9595-8.

Daoust, Marc-Kevin. 2019. “One Standard To Rule Them All?” *Ratio* 32 (1): 12-21.

Daoust, Marc-Kevin. “The Deontic Force of Apparent Reasons.” Unpublished manuscript a

Daoust, Marc-Kevin. “A Puzzle About the Optimization of Collective Reliability.” Unpublished manuscript b

Daoust, Marc-Kevin. “A Note on Accuracy-Dominance Vindications of Consistency.” Unpublished manuscript c

Daoust, Marc-Kevin. “Should Agents Be Immodest?” Unpublished manuscript d

Daoust, Marc-Kevin. “What Second-Best Epistemology Could Look Like” Unpublished manuscript e

Daoust, Marc-Kevin, and David Montminy. 2017. “Testing Minimax for Rational Ignorant Agents”: 541-553, in *Logic, Rationality, and Interaction*, edited by A. Baltag *et al.*, LNCS 10455, Berlin: Springer

Daoust, Marc-Kevin, and David Montminy. “Immodesty and Permissivism.” Unpublished manuscript

Davidson, Donald. 1982. “Paradoxes of Irrationality.” In *Philosophical Essays on Freud*, edited by Richard Wollheim and James Hopkins, 289–305. Cambridge: Cambridge University Press.

De Bona, Glauber, and Julia Staffel. 2018. “Why Be (Approximately) Coherent?” *Analysis*, 1–11. https://doi.org/10.1093/analys/anx159.

Demey, Lorenz. 2013. “Contemporary Epistemic Logic and the Lockean Thesis.” *Foundations of Science* 18 (4): 599–610.

Dietrich, Franz, and Christian List. 2004. “A Model of Jury Decisions Where All Jurors Have the Same Evidence.” *Synthese* 142 (2): 175–202.

Dietrich, Franz, and Kai Spiekermann. 2013. “Epistemic Democracy With Defensible Premises.” *Economics & Philosophy* 29 (1): 87–120.

DiPaolo, Joshua. 2018. “Second Best Epistemology: Fallibility and Normativity.” *Philosophical Studies*, 1–24. https://doi.org/10.1007/s11098-018-1110-y.

Dogramaci, Sinan, and Sophie Horowitz. 2016. “An Argument for Uniqueness About Evidential Support.” *Philosophical Issues* 26 (1): 130–47.

Dorst, Kevin. 2017. “Lockeans Maximize Expected Accuracy.” *Mind*. https://doi.org/10.1093/mind/fzx028.

Douven, Igor. 2009. “Uniqueness Revisited.” *American Philosophical Quarterly*, 347–361.

Drake, J. 2016. “Doxastic Permissiveness and the Promise of Truth.” *Synthese*, 1–16. https://doi.org/10.1007/s11229-016-1176-z.

Easwaran, Kenny. 2015. “Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities.” *Noûs*. https://doi.org/10.1111/nous.12099/pdf.

Easwaran, Kenny, and Branden Fitelson. 2015. “Accuracy, Coherence, and Evidence.” *Oxford Studies in Epistemology* 5: 61–96.

Elga, Adam. 2010. “How to Disagree about How to Disagree.” In *Disagreement*, edited by Ted Warfield and Richard Feldman. Oxford: Oxford University Press.

Elga, Adam. 2013. “The Puzzle of the Unmarked Clock and the New Rational Reflection Principle.” *Philosophical Studies* 164 (1): 127–39. https://doi.org/10.1007/s11098-013-0091-0.

Engel, Pascal. 2004. “Truth and the Aim of Belief.” In *Laws and Models in Science*, edited by D. Gillies, 77–97. London: King’s College Publications.

Estlund, David. 1994. “Opinion Leaders, Independence, and Condorcet’s Jury Theorem.” *Theory and Decision* 36 (2): 131–162.

Estlund, David. 2014. “Utopophobia.” *Philosophy & Public Affairs* 42 (2): 113–134.

Feldman, Richard. 1988. “Having Evidence.” In *Philosophical Analysis: A Defense by Example*, edited by David F. Austin, 83–104. Philosophical Studies Series. Dordrecht: Springer Netherlands.

Feldman, Richard. 2005a. “Justification Is Internal.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 337–50. John Wiley & Sons.

Feldman, Richard. 2005b. “Respecting the Evidence.” *Philosophical Perspectives* 19 (1): 95–119.

Fey, Mark. 2003. “A Note on the Condorcet Jury Theorem with Supermajority Voting Rules.” *Social Choice and Welfare* 20 (1): 27–32.

Fitelson, Branden. m.s. *Coherence*.

Fogal, Daniel. m.s. “Structural Requirements and the Primacy of Pressure.” *Unpublished Manuscript*.

Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge (MA): Harvard University Press.

Foley, Richard. 1993. *Working without a Net*. Oxford: Oxford: Oxford University Press.

Foley, Richard. 2009. "Beliefs, Degrees of Belief, and the Lockean Thesis." In *Degrees of Belief*, 37–47. Springer.

Gaus, Gerald. 2016. *The Tyranny of the Ideal: Justice in a Diverse Society*. Princeton and Oxford: Princeton University Press.

Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford: Oxford University Press.

Gibbons, John. 2013. *The Norm of Belief*. Oxford: Oxford University Press.

Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.

Goldman, Alvin. 2010. "Epistemic Relativism and Reasonable Disagreement." In *Disagreement*, edited by Ted Warfield and Richard Feldman, 187–213. Oxford: Oxford University Press.

Goldman, Alvin. 2015. "Reliabilism, Veritism, and Epistemic Consequentialism." *Episteme* 12 (02): 131–143.

Greco, Daniel. 2014. "A Puzzle about Epistemic Akrasia." *Philosophical Studies* 167 (2): 201–219.

Greco, Daniel, and Brian Hedden. 2016. "Uniqueness and Metaepistemology." *The Journal of Philosophy* 113 (8): 365–95. https://doi.org/10.5840/jphil2016113825.

Greco, John. 2005. "Justification Is Not Internal." In *Contemporary Debates in Epistemology*. Oxford: Blackwell Publishers.

Grimm, Stephen R. 2008. "Epistemic Goals and Epistemic Values." *Philosophy and Phenomenological Research* 77 (3): 725–744.

Grimm, Stephen R. 2009. "Epistemic Normativity." In *Epistemic Value*, edited by Adrian Haddock, Allan Millar, and Duncan Pritchard, 2009:243–64. Oxford, New York: Oxford University Press.

Guindon, Bruno. 2014. "Sources, Raisons et Exigences." *Les Ateliers de l'éthique/The Ethics Forum* 9: 152–165.

Guindon, Bruno. 2016. "Sources, Reasons, and Requirements." *Philosophical Studies* 173 (5): 1253–1268.

Harman, Gilbert. 1999. *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press.

Hatcher, Michael. 2016. "Accessibilism Defined." *Episteme*, 1–23. https://doi.org/10.1017/epi.2016.36.

Hawthorne, James. 2009. "The Lockean Thesis and the Logic of Belief." In *Degrees of Belief*, 49–74. Springer.

Hedden, Brian. 2015a. "Time-Slice Rationality." *Mind* 124 (494): 449–491.

Hedden, Brian. 2015b. "A Defense of Objectivism about Evidential Support." *Canadian Journal of Philosophy* 45 (5–6): 716–743.

Hieronymi, Pamela. 2005. “The Wrong Kind of Reason.” *The Journal of Philosophy* 102 (9): 437–457.

Hinchman, Edward S. 2013. “Rational Requirements and ‘Rational’ Akrasia.” *Philosophical Studies* 166 (3): 529–552.

Hlobil, Ulf. 2015. “There Are Diachronic Norms of Rationality.” *Thought: A Journal of Philosophy* 4 (1): 38–45.

Hong, Lu, and Scott Page. 2012. “The Micro-Foundations of Collective Wisdom.” In *Collective Wisdom: Principles and Mechanisms*, edited by Jon Elster and Hélène Landemore, 56–71. Cambridge: Cambridge University Press.

Horowitz, Sophie. 2014a. “Epistemic Akrasia.” *Noûs* 48 (4): 718–744.

Horowitz, Sophie. 2014b. “Immoderately Rational.” *Philosophical Studies* 167 (1): 41–56.

Horwich, Paul. 2010. *Truth, Meaning, Reality*. Oxford, New York: Oxford University Press.

Horwich, Paul. 2013. “Belief-Truth Norms.” In *The Aim of Belief*, edited by Timothy Chan. Oxford: Oxford University Press.

Huemer, Michael. 2007. “Moore’s Paradox and the Norm of Belief.” In *Themes from GE Moore*, edited by Susana Nucceteli and Gary Seays, 142–157. New York: Oxford University Press.

Hughes, Nick. 2017. “Dilemmic Epistemology.” *Synthese*. https://doi.org/10.1007/s11229-017-1639-x.

Joyce, James M. 1998. “A Nonpragmatic Vindication of Probabilism.” *Philosophy of Science* 65 (4): 575–603.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica: Journal of the Econometric Society* 47 (2): 263–291.

Kaniovski, Serguei. 2010. “Aggregation of Correlated Votes and Condorcet’s Jury Theorem.” *Theory and Decision* 69 (3): 453–468.

Kearns, Stephen, and Daniel Star. 2008. “Reasons: Explanations or Evidence?” *Ethics* 119 (1): 31–56.

Kelly, Thomas. 2005. “The Epistemic Significance of Disagreement.” In *Oxford Studies in Epistemology*, edited by Tamar Szabo´ Gendler and John Hawthorne, 1:167–96. Oxford: Clarendon Press.

Kelly, Thomas. 2014. “Evidence Can Be Permissive.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 298–312. John Wiley & Sons.

Kiesewetter, Benjamin. 2012. “A Dilemma for Parfit’s Conception of Normativity.” *Analysis* 72 (3): 466–474.

Kiesewetter, Benjamin. 2013. *The Normativity of Rationality*. Humboldt-Universität zu Berlin: Dissertation.

Kiesewetter, Benjamin. 2016. “You Ought to ϕ Only If You May Believe That You Ought to ϕ.” *The Philosophical Quarterly* 66 (265): 760–82. https://doi.org/10.1093/pq/pqw012.

Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.

Kiesewetter, Benjamin. 2018. “Can The Lottery Paradox be Solved by Identifying Epistemic Justification with Epistemic Permissibility?” *Episteme*, 1-21. https://doi.org/10.1017/epi.2018.4

Kolodny, Niko. 2005. “Why Be Rational?” *Mind* 114 (455): 509–563.

Kolodny, Niko. 2007a. “State or Process Requirements?” *Mind* 116 (462): 371–385.

Kolodny, Niko. 2007b. “How Does Coherence Matter?” In *Proceedings of the Aristotelian Society*, 229–263.

Kolodny, Niko. 2008a. “The Myth of Practical Consistency.” *European Journal of Philosophy* 16 (3): 366–402.

Kolodny, Niko. 2008b. “Why Be Disposed to Be Coherent?” *Ethics* 118 (3): 437–463.

Kopec, Matthew. 2015. “A Counterexample to the Uniqueness Thesis.” *Philosophia* 43 (2): 403–409.

Kopec, Matthew, and Michael G. Titelbaum. 2016. “The Uniqueness Thesis.” *Philosophy Compass* 11 (4): 189–200.

Kroedel, Thomas. 2011. “The Lottery Paradox, Epistemic Justification and Permissibility.” *Analysis* 72 (1): 57–60.

Kroedel, T. 2013a. “The Permissibility Solution to the Lottery Paradox – Reply to Littlejohn.” *Logos & Episteme*, 4: 103–11.

Kroedel, T. 2013b. “Why Epistemic Permissions Don’t Agglomerate – Another Reply to Littlejohn.” *Logos & Episteme*, 4: 451–5.

Kroedel, T. 2017. “The Lottery, the Preface, and Conditions on Permissible Belief.” *Erkenntnis*. https://doi.org/10.1007/s10670-017-9911-5.

Kvanvig, Jonathan. 2008. “Pointless Truth.” *Midwest Studies in Philosophy* 32 (1): 199–212.

Kvanvig, Jonathan. 2009. “The Value of Understanding.” In *Epistemic Value*, edited by Adrian Haddock, Allan Millar, and Duncan Pritchard, 95–111. Oxford, New York: Oxford University Press.

Landemore, Hélène. 2012. “Democratic Reason: The Mechanisms of Collective Intelligence in Politics.” In *Collective Wisdom: Principles and Mechanisms*, edited by Hélène Landemore and Jon Elster, 251–89. Cambridge: Cambridge University Press.

Landemore, Hélène. 2013. “Deliberation, Cognitive Diversity, and Democratic Inclusiveness: An Epistemic Argument for the Random Selection of Representatives.” *Synthese* 190 (7): 1209–31. https://doi.org/10.1007/s11229-012-0062-6.

Lasonen-Aarnio, Maria. “Enkrasia or Evidentialism? Learning to Love Mismatch.” *Philosophical Studies*, Forthcoming.

Lasonen-Aarnio, Maria. 2014. “Higher-Order Evidence and the Limits of Defeat.” *Philosophy and Phenomenological Research* 88 (2): 314–345.

Lasonen-Aarnio, Maria. 2015. “New Rational Reflection and Internalism about Rationality.” *Oxford Studies in Epistemology* 5: 145–171.

Lehrer, Keith. 1990. *Theory of Knowledge*. Routledge.

Leitgeb, Hannes. 2014a. “The Review Paradox: On the Diachronic Costs of Not Closing Rational Belief under Conjunction.” *Noûs* 48 (4): 781–793.

Leitgeb, Hannes. 2014b. “The Stability Theory of Belief.” *Philosophical Review* 123 (2): 131–171.

Leitgeb, Hannes, and Richard Pettigrew. 2010. “An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy*.” *Philosophy of Science* 77 (2): 236–272.

Levinstein, Benjamin Anders. 2017. “A Pragmatist’s Guide to Epistemic Utility.” *Philosophy of Science* 84 (4): 613–638.

Lewis, David. 1971. “Immodest Inductive Methods.” *Philosophy of Science* 38 (1): 54–63.

Li, Han. 2017. “A Theory of Epistemic Supererogation.” *Erkenntnis*, 1–19. https://doi.org/10.1007/s10670-017-9893-3.

Lipsey, Richard G., and Kelvin Lancaster. 1956. “The General Theory of Second Best.” *The Review of Economic Studies* 24 (1): 11–32.

List, Christian, and Robert E. Goodin. 2001. “Epistemic Democracy: Generalizing the Condorcet Jury Theorem.” *Journal of Political Philosophy* 9 (3): 277–306.

Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.

Littlejohn, Clayton. 2015. “Stop Making Sense? On a Puzzle about Rationality.” *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12271.

Lockard, Matthew. 2013. “Epistemic Instrumentalism.” *Synthese* 190 (9): 1701–18.

Loftus, Elizabeth F. 2005. “Planting Misinformation in the Human Mind: A 30-Year Investigation of the Malleability of Memory.” *Learning & Memory* 12 (4): 361–366.

Longino, Helen. 2016. “The Social Dimensions of Scientific Knowledge.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2016. http://plato.stanford.edu/archives/spr2016/entries/scientific-knowledge-social/.

Lord, Errol. 2010. “Having Reasons and the Factoring Account.” *Philosophical Studies* 149 (3): 283–296.

Lord, Errol. 2014. “The Coherent and the Rational.” *Analytic Philosophy* 55 (2): 151–175.

Lord, Errol. 2015. “Acting for the Right Reasons, Abilities, and Obligation.” In *Oxford Studies in Metaethics*, edited by Russ Shafer-Landau, 10:26–52. Oxford: Oxford University Press.

Lord, Errol. 2017. “What You’re Rationally Required to Do and What You Ought to Do (Are the Same Thing!).” *Mind*. https://doi.org/10.1093/mind/fzw023.

Lord, Errol. 2018. *The Importance of Being Rational*. Oxford: Oxford University Press.

Lynch, Michael. 2004. *True to Life: Why Truth Matters*. Cambridge (MA): MIT Press.

Lynch, Michael. 2009. “The Truth of Values and the Values of Truth.” In *Epistemic Value*, edited by Adrian Haddock, Allan Millar, and Duncan Pritchard. Oxford, New York: Oxford University Press.

Matheson, Jonathan. 2011. “The Case for Rational Uniqueness.” *Logos & Episteme* 2 (3): 359–373.

Mayo-Wilson, Conor, and Gregory Wheeler. 2016. “Scoring Imprecise Credences: A Mildly Immodest Proposal.” *Philosophy and Phenomenological Research*. http://onlinelibrary.wiley.com/doi/10.1111/phpr.12256/full.

McCain, Kevin. 2008. “The Virtues of Epistemic Conservatism.” *Synthese* 164 (2): 185–200.

McGrath, Matthew. 2007. “Memory and Epistemic Conservatism.” *Synthese* 157 (1): 1–24.

McGrath, Matthew, and Jeremy Fantl. 2013. “Practical Matters Affect Whether You Know.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 84–95. John Wiley & Sons.

McHugh, Conor. 2012. “The Truth Norm of Belief.” *Pacific Philosophical Quarterly* 93 (1): 8–30.

McHugh, Conor, and Jonathan Way. 2017. “Objectivism and Perspectivism about the Epistemic Ought.” *Ergo* 4 (5): 121–45. https://doi.org/10.3998/ergo.12405314.0004.005.

Meacham, Christopher JG. 2014. “Impermissive Bayesianism.” *Erkenntnis* 79 (6): 1185–1217.

Mele, Alfred. 1988. *Irrationality*. Oxford: Oxford University Press.

Mitova, Veli. 2017. *Believable Evidence*. Cambridge: Cambridge University Press.

Nagel, Thomas. 1978. *The Possibility of Altruism*. Princeton University Press.

Olinder, Ragnar Francén. 2012. “Rescuing Doxastic Normativism.” *Theoria* 78 (4): 293–308. https://doi.org/10.1111/j.1755-2567.2012.01144.x.

Owens, David. 2002. *Reason Without Freedom: The Problem of Epistemic Normativity*. Routledge.

Ozga, S. A. 1955. “An Essay in the Theory of Tariffs.” *Journal of Political Economy* 63 (6): 489–99. https://doi.org/10.1086/257723.

Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton (NJ): Princeton University Press.

Page, Scott E. 2010. *Diversity and Complexity*. Princeton (NJ): Princeton University Press.

Parfit, Derek. 2011. *On What Matters: Volume One*. Oxford University Press.

Pasnau, Robert. 2013. "Epistemology Idealized." *Mind* 122 (488): 987–1021. https://doi.org/10.1093/mind/fzt093.

Paul, L. A., and John Quiggin. 2018. "Real World Problems." *Episteme*, 1–20. https://doi.org/10.1017/epi.2018.28.

Pears, David Francis. 1984. *Motivated Irrationality*. Oxford: Oxford University Press.

Peijnenburg, Jeanne, and David Atkinson. 2011. "Grounds and Limits: Reichenbach and Foundationalist Epistemology." *Synthese* 181 (1): 113–124.

Pettigrew, Richard. 2016a. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Pettigrew, Richard. 2016b. "Accuracy, Risk, and the Principle of Indifference." *Philosophy and Phenomenological Research* 92 (1): 35–59. https://doi.org/10.1111/phpr.12097.

Pettigrew, Richard. 2016c. "Jamesian Epistemology Formalised: An Explication Of 'The Will To Believe.'" *Episteme* 13 (3): 253–68. https://doi.org/10.1017/epi.2015.44.

Podgorski, Abelard. 2016a. "A Reply to the Synchronist." *Mind* 125 (499): 859–871.

Podgorski, Abelard. 2016b. "Dynamic Conservatism." *Ergo, an Open Access Journal of Philosophy* 3. https://doi.org/10.3998/ergo.12405314.0003.013.

Podgorski, Abelard. 2016c. "Dynamic Permissivism." *Philosophical Studies* 173 (7): 1923–1939.

Polikar, Robi. 2012. "Ensemble Learning." In *Ensemble Machine Learning: Methods and Applications*, edited by C Zhang and Y MA, 1–34. Springer.

Prichard, Harold Arthur. 1912. "Does Moral Philosophy Rest on a Mistake?" *Mind* 21 (81): 21–37.

Pritchard, Duncan. 2016 *Epistemic Angst*. Princeton: Princeton University Press

Pryor, James. 2013. "Problems for Credulism." *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, 89–131.

Railton, Peter. 2014. "The Affective Dog and Its Rational Tale: Intuition and Attunement." *Ethics* 124 (4): 813–859.

Raleigh, Thomas. 2013. "Belief Norms and Blindspots." *The Southern Journal of Philosophy* 51 (2): 243–269.

Raleigh, Thomas. 2015. "An Argument for Permissivism from Safespots." In *Logic, Rationality, and Interaction*, edited by Wiebe van der Hoek, Wesley H. Holliday, and

Wen-fang Wang, 308–15. Lecture Notes in Computer Science 9394. Berlin: Springer. https://doi.org/10.1007/978-3-662-48561-3_25.

Reisner, Andrew. 2007. “Evidentialism and the Numbers Game.” *Theoria* 73 (4): 304–316.

Reisner, Andrew. 2009a. “The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem.” *Philosophical Studies* 145 (2): 257–272.

Reisner, Andrew. 2009b. “Unifying the Requirements of Rationality.” *Philosophical Explorations* 12 (3): 243–60. https://doi.org/10.1080/13869790903067709.

Reisner, Andrew. 2011. “Is There Reason to Be Theoretically Rational?” In *Reasons for Belief*, edited by Andrew Reisner and Asbjørn Steglich-Petersen. Cambridge: Cambridge University Press.

Reisner, Andrew. 2013. “Is the Enkratic Principle a Requirement of Rationality?” *Organon F.* 20 (4): 437–63.

Reisner, Andrew. 2014. “A Short Refutation of Strict Normative Evidentialism.” *Inquiry*, 1–9. https://doi.org/10.1080/0020174X.2014.932303.

Reisner, Andrew. 2015. “Pragmatic Reasons for Belief.” *Forthcoming*. http://philpapers.org/rec/REIPRF.

Raz, Joseph. 1999. “Explaining Normativity: On Rationality and the Justification of Reason.” *Ratio* 12 (4): 354–379.

Raz, Joseph. 2005. “The Myth of Instrumental Rationality.” *J. Ethics & Soc. Phil.* 1: 1–28.

Ribeiro, Brian. 2011. “Epistemic Akrasia.” *International Journal for the Study of Skepticism* 1 (1): 18–25.

Rokach, Lior. 2010. “Ensemble-Based Classifiers.” *Artificial Intelligence Review* 33 (1): 1–39.

Romeijn, Jan-Willem, and David Atkinson. 2011. “Learning Juror Competence: A Generalized Condorcet Jury Theorem.” *Politics, Philosophy & Economics* 10 (3): 237–262.

Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge (MA): Harvard University Press.

Schellenberg, Susanna. 2016. “Phenomenal Evidence and Factive Evidence.” *Philosophical Studies* 173: 875–96.

Schoenfield, Miriam. 2014. “Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief.” *Noûs* 48 (2): 193–218.

Schroeder, Mark. 2008. “Having Reasons.” *Philosophical Studies* 139 (1): 57–71.

Schroeder, Mark. 2011. “What Does It Take to ‘Have’ a Reason.” In *Reasons for Belief*, edited by Andrew Reisner and Asbjørn Steglich-Petersen, 201–222. Cambridge: Cambridge University Press.

Schroeder, Mark. 2015. “What Makes Reasons Sufficient?” *American Philosophical Quarterly* 52 (2): 159–70.

Shah, Nishi. 2003. “How Truth Governs Belief.” *The Philosophical Review* 112 (4): 447–482.

Shah, Nishi, and J. David Velleman. 2005. “Doxastic Deliberation.” *The Philosophical Review* 114 (4): 497–534.

Sharadin, Nathaniel. 2015. “A Partial Defense of Permissivism.” *Ratio*. https://doi.org/10.1111/rati.12115.

Shpall, Sam. 2013. “Wide and Narrow Scope.” *Philosophical Studies* 163 (3): 717–736.

Siegel, Harvey. 2019. “Hinges, Disagreements, and Arguments: (Rationally) Believing Hinge Propositions and Arguing across Deep Disagreements.” *Topoi*. https://doi.org/10.1007/s11245-018-9625-6

Sinnott-Armstrong, Walter. 1996. “Moral Dilemmas and Rights.” In *Moral Dilemmas and Moral Theory*, edited by H.E. Mason. New York: Oxford University Press.

Sinnott-Armstrong, Walter. 2012. “Moore’s Paradox and the Accessibility of Justification.” *Philosophy and Phenomenological Research* 85 (2): 273–300.

Smithies, Declan. 2015. “Ideal Rationality and Logical Omniscience.” *Synthese* 192 (9): 2769–93. https://doi.org/10.1007/s11229-015-0735-z.

Smithies, Declan. 2016. “Belief and Self-Knowledge: Lessons From Moore’s Paradox.” *Philosophical Issues* 26 (1): 393–421. https://doi.org/10.1111/phis.12075.

Snedegar, Justin. 2018. “Reasons for and Reasons Against.” *Philosophical Studies* 175 (3): 725–743.

Southwood, Nicholas. 2008. “Vindicating the Normativity of Rationality.” *Ethics* 119 (1): 9–30.

Staffel, Julia. 2017. “Should I Pretend I’m Perfect?” *Res Philosophica* 94 (2): 301–24. https://doi.org/10.11612/resphil.1536.

Steinberger, Florian. Forthcoming. “Accuracy and Epistemic Conservatism” *Analysis*. https://philpapers.org/rec/STEAAE-4

Stemplowska, Zofia, and Adam Swift. 2012. “Ideal and Nonideal Theory.” In *The Oxford Handbook of Political Philosophy*, edited by David Estlund, 373–90. Oxford: OUP.

Stone, Peter. 2015. “Introducing Difference into the Condorcet Jury Theorem.” *Theory and Decision* 78 (3): 399–409.

Sturgeon, Scott. 2008. “Reason and the Grain of Belief.” *Noûs* 42 (1): 139–165.

Sunstein, Cass R. 2006. *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press.

Surowiecki, James. 2005. *The Wisdom of Crowds*. New York: Anchor.

Sylvan, Kurt. m.s. “On Divorcing the Rational and the Justified.” *Unpublished Manuscript*.

Sylvan, Kurt. 2015a. “What Apparent Reasons Appear to Be.” *Philosophical Studies* 172 (3): 587–606. https://doi.org/10.1007/s11098-014-0320-1.

Sylvan, Kurt. 2015b. “The Illusion of Discretion.” *Synthese*, 1–31. https://doi.org/10.1007/s11229-015-0796-z.

Sylvan, Kurt. 2016a. “Epistemic Reasons I: Normativity.” *Philosophy Compass* 11 (7): 364–76. https://doi.org/10.1111/phc3.12327.

Sylvan, Kurt. 2016b. “Epistemic Reasons II: Basing.” *Philosophy Compass* 11 (7): 377–89. https://doi.org/10.1111/phc3.12329.

Talbott, William. 2016. “Bayesian Epistemology.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian/.

Tappolet, Christine. 2014. “The Normativity of Evaluative Concepts.” In *Mind, Values, and Metaphysics*, edited by Anne Reboul, 39–54. Springer.

Titelbaum, Michael G. Forthcoming. “Return to Reason.” In *Higher-Order Evidence: New Essays*, edited by Matthias Skipper and Asbjørn Steglich-Petersen. Oxford: Oxford University Press.

Titelbaum, Michael G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford: Oxford University Press.

Titelbaum, Michael G. 2015a. “How to Derive a Narrow-Scope Requirement from Wide-Scope Requirements.” *Philosophical Studies* 172 (2): 535–542.

Titelbaum, Michael G. 2015b. “Rationality’s Fixed Point (or: In Defense of Right Reason).” *Oxford Studies in Epistemology* 5: 253–294.

Titelbaum, Michael G., and Matthew Kopec. m.s. “Plausible Permissivism.” *Unpublished Manuscript*.

Titelbaum, Michael G., and Matthew Kopec. Forthcoming. “When Rational Reasoners Reason Differently.” In *Reasoning: Essays on Theoretical and Practical Thinking*, edited by M Balcerak Jackson and B Balcerak Jackson. Oxford University Press.

Todd, Peter M., and Gerd Gigerenzer. 2000. “Précis of Simple Heuristics That Make Us Smart.” *Behavioral and Brain Sciences* 23 (05): 727–741.

Turri, John. 2009. “The Ontology of Epistemic Reasons.” *Noûs* 43 (3): 490–512.

Vahid, Hamid. 2004. “Varieties of Epistemic Conservatism.” *Synthese* 141 (1): 97–122.

Vermeule, Adrian. 2012. “Collective Wisdom and Institutional Design.” In *Collective Wisdom: Principles and Mechanisms*, edited by Hélène Landemore and Jon Elster, 338–67. Cambridge: Cambridge University Press.

Vineberg, Susan. 2016. “Dutch Book Arguments.” Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/spr2016/entries/dutch-book/.

Walton, Douglas. 1988. “Burden of Proof.” *Argumentation* 2: 233–54.

Way, Jonathan. 2009. “Two Accounts of the Normativity of Rationality.” *Journal of Ethics and Social Philosophy* 4(1): 1-9.

Way, Jonathan. 2010a. “Defending the Wide-Scope Approach to Instrumental Reason.” *Philosophical Studies* 147 (2): 213.

Way, Jonathan. 2010b. “The Normativity of Rationality.” *Philosophy Compass* 5 (12): 1057–1068.

Way, Jonathan. 2011. “The Symmetry of Rational Requirements.” *Philosophical Studies* 155 (2): 227–239.

Way, Jonathan. 2012. “Explaining the Instrumental Principle.” *Australasian Journal of Philosophy* 90 (3): 487–506.

Way, Jonathan, and Daniel Whiting. 2016a. “Reasons and Guidance.” *Analytic Philosophy* 57 (3): 214–35.

Way, Jonathan, and Daniel Whiting. 2016b. “If You Justifiably Believe That You Ought to Φ, You Ought to Φ.” *Philosophical Studies* 173: 1873–95. https://doi.org/10.1007/s11098-015-0582-2.

Weatherson, Brian. 2015. “Memory, Belief and Time.” *Canadian Journal of Philosophy* 45 (5–6): 692–715.

Wedgwood, Ralph. 2002. “Internalism Explained*.” *Philosophy and Phenomenological Research* 65 (2): 349–69. https://doi.org/10.1111/j.1933-1592.2002.tb00206.x.

Wedgwood, Ralph. 2013. “Right Thing to Believe.” In *The Aim of Belief*, edited by Timothy Chan. Oxford: Oxford University Press.

Wedgwood, Ralph. 2015. “The Pitfalls of ‘Reasons.’” *Philosophical Issues* 25 (1): 123–143.

Wedgwood, Ralph. 2017. *The Value of Rationality*. Oxford: Oxford University Press.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich. 2001. “Normativity and Epistemic Intuitions.” *Philosophical Topics* 29 (1/2): 429–460.

Weisberg, Jonathan. 2015. “Updating, Undermining, and Independence.” *The British Journal for the Philosophy of Science* 66 (1): 121–59.

White, Roger. 2005. “Epistemic Permissiveness.” *Philosophical Perspectives* 19 (1): 445–459.

White, Roger. 2014. “Evidence Cannot Be Permissive.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, and Ernest Sosa, 312–23. John Wiley & Sons.

Whiting, Daniel. 2010. “Should I Believe the Truth?” *Dialectica* 64 (2): 213–24. https://doi.org/10.1111/j.1746-8361.2009.01204.x.

Whiting, Daniel. 2013. “The Good and the True (or the Bad and the False).” *Philosophy* 88 (02): 219–242.

Williams, Bernard AO. 1965. “Ethical Consistency.” *Proceedings of the Aristotelian Society, Supplementary Volume 35*, 103–124.

Williamson, Jon. 2010. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.

Williamson, Timothy. 2002. *Knowledge and Its Limits*. Oxford University Press.

Williamson, Timothy. 2008. *The Philosophy of Philosophy*. John Wiley & Sons.

Worsnip, Alex. Forthcoming. “Isolating Correct Reasoning.” In *Reasoning: New Essays on Theoretical and Practical Thinking*, edited by Magdalena Balcerak Jackson and Brendan Balcerak Jackson.

Worsnip, Alex. m.s. “Disagreement as Interpersonal Incoherence.” Unpublished manuscript

Worsnip, Alex. 2015. “The Conflict of Evidence and Coherence.” *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12246.

Worsnip, Alex. 2016. “Moral Reasons, Epistemic Reasons and Rationality.” *The Philosophical Quarterly* 66 (263): 341–361.

Worsnip, Alex. 2018. “Reasons, Rationality, Reasoning: How Much Pulling Apart?” *Problema* 12: 59–93. https://doi.org/10.22201/iij.24487937e.2018.12.12443.