

Dispositional Theories of Knowledge

A Defence of Aetiological Foundationalism

Lars Bo Gundersen

Dispositional Theories of Knowledge

A Defence of Aetiological Foundationalism

LARS BO GUNDERSEN

Aarhus University, Denmark

ASHGATE

© Lars Bo Gundersen 2003

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior permission of the publisher.

Lars Bo Gundersen has asserted his moral right under the Copyright, Designs and Patents Act, 1988, to be identified as the author of this work.

Published by
Ashgate Publishing Limited
Gower House
Croft Road
Aldershot
Hampshire GU11 3HR
England

Ashgate Publishing Company
Suite 420
101 Cherry Street
Burlington, VT 05401-4405
USA

Ashgate website: http://www.ashgate.com
--

British Library Cataloguing in Publication Data

Gundersen, Lars Bo

Dispositional theories of knowledge : a defence of
aetiological foundationalism. – (Ashgate new critical
thinking in philosophy)

1. Knowledge, Theory of

I. Title

121

Library of Congress Cataloging-in-Publication Data

Gundersen, Lars Bo.

Dispositional theories of knowledge : defence of aetiological foundationalism /
Lars Bo Gundersen.

p. cm. – (Ashgate new critical thinking in philosophy)

Includes bibliographical references and index.

ISBN 0-7546-3051-X (hardback)

1. Knowledge, Theory of. 2. Disposition (Philosophy) 3. Modality (Theory of
knowledge) I. Title. II. Series.

BD161 .G86 2002

121–dc21

2002027699

ISBN 0 7546 3051 X

Printed and bound in Great Britain by MPG Books Ltd, Bodmin, Cornwall



Master Nakano practising Kyudo ('The Way of the Bow'). From *Martial Arts*, Peter Payne, Thames and Hudson (1981). Reproduced here by permission of Michel Random.

DISPOSITIONAL THEORIES OF KNOWLEDGE

This book offers an original examination of human cognition, arguing that cognitive skills are dispositional in nature. Opposing influential views in modern Anglo-American philosophy, Gundersen starts from the received premis that knowledge is analyzable in terms of belief, justification and truth, and goes on to clarify and improve on these ingredients' exact nature and internal association. Exploring a wide range of arguments offered by influential contributors in the field of modal epistemology, Gundersen argues that external conditions are secondary in developing and cultivating cognitive competence and that the fulcrum of the cognitive investigation is the fascinating interplay between and cultivation of internal cognitive powers.

ASHGATE NEW CRITICAL THINKING IN PHILOSOPHY

The *Ashgate New Critical Thinking in Philosophy* series aims to bring high quality research monograph publishing back into focus for authors, the international library market, and student, academic and research readers. Headed by an international editorial advisory board of acclaimed scholars from across the philosophical spectrum, this new monograph series presents cutting-edge research from established as well as exciting new authors in the field; spans the breadth of philosophy and related disciplinary and interdisciplinary perspectives; and takes contemporary philosophical research into new directions and debate.

Series Editorial Board:

David Cooper, University of Durham, UK
Peter Lipton, University of Cambridge, UK
Sean Sayers, University of Kent at Canterbury, UK
Simon Critchley, University of Essex, UK
Simon Glendinning, University of Reading, UK
Paul Helm, King's College London, UK
David Lamb, University of Birmingham, UK
John Post, Vanderbilt University, Nashville, USA
Alan Goldman, University of Miami, Florida, USA
Joseph Friggieri, University of Malta, Malta
Graham Priest, University of Melbourne, Australia
Moirá Gatens, University of Sydney, Australia
Alan Musgrave, University of Otago, New Zealand

For Andy Law

Acknowledgements

The present work has benefited from stimulating discussions with a wide range of philosophers. In particular, I would like to thank Lars Binderup, Eline Busck, Jonathan Dancy, Patrick Greenough, Bob Hale, Jesper Kallestrup, Fraser McBride, Duncan Pritchard, Stephen Read, Sven Rosenkranz and Stephen Shapiro for their valuable criticism and suggestions for amendments. I would also like to thank the students at my upper-level courses at the philosophical departments at Copenhagen University and the University of Aarhus to whom I had the privileged opportunity to communicate some of the ideas contained here. I learned a lot from the lively discussions succeeding each seminar. I would like to express a very special gratitude to Crispin Wright who has influenced my thinking during the writing process to such an extent that I cannot clearly distinguish any more between the ideas that originated from conversations with him and those that arose from my own thinking. Thus, I attribute to him whatever share of the ideas that he, or anyone else, might wish. Finally I would like to express my gratitude to my editor, Sarah Charters for her efficiency and kindness in helping to prepare the final manuscript, and to Jens Kristian Mathiasen for invaluable assistance with the index.

Contents

<i>Acknowledgements</i>	ix
1 The General Idea Underlying Modal Epistemology	1
2 Modal Theories of Knowledge and Warranted Belief	12
3 Tracking Dispositionality	48
4 Masking, Mimicking and Altering	69
5 A Dispositional Theory of Knowledge and Warranted Belief	106
<i>Bibliography</i>	143
<i>Index</i>	148

Chapter 1

The General Idea Underlying Modal Epistemology

Introduction

Once upon a time it seemed a simple matter to define knowledge: an agent S knew that P, the contention then was, if and only if S had a justified true belief that P. But then Gettier came along and reminded everyone that justified true beliefs do not always suffice for knowledge.¹ Gettier-type considerations can be illustrated by the following example: the Clock case. If Foucault, whose clock happened to stop exactly 24 hours ago, forms a belief as to what time it is, then, although this belief is justified (by looking at his clock) and true (the clock actually shows, just now, what time it is), it by no means counts as knowledge. Foucault is extremely lucky that he gets the time right in this instance. Had he consulted his stopped clock 5 minutes earlier, or later, he would have got the time wrong. And knowledge, properly understood, does not allow for such an amount of good fortune.

Gettier's reminder that knowledge cannot be identified with justified true belief kicked off a heated debate about what knowledge *really* is, and, further, whether knowledge allows for conceptual analysis at all. In the course of this debate it has been suggested that each of the ingredient conditions in the classical tripartite definition—belief, justification and truth—calls for revision, both in order for these ingredients jointly to suffice for knowledge, but also, and maybe more surprisingly, for each individually to count as necessary. The motivation driving the present inquiry is that some conceptually illuminating stories yet remain to be told. I will assume, as a starting point, that knowledge is related, somehow, to belief, justification² and truth, and regard the challenge as to clarify and improve on these ingredients' exact nature and internal association.

¹ Gettier 1963. Gettier was not the first to publish counterexamples to the classical tripartite definition of knowledge. Lewis Carroll's *A Tangled Tale* contains suitable candidates (thanks to James Chase for pointing this out). Gettier was not even the first to publish Gettier counterexamples in a strictly philosophical context. The clock case appears in Russell's *Human Knowledge*, published in 1948 (p. 113).

² Historically, the term 'justification' has been associated with internalist accounts of justification, that by which the cognitive agent can justify his belief in P if challenged to do so. There are, however, no a priori grounds for thinking that any such justificatory feature of the epistemic situation has to be internal to the agent. The use of the term 'justification' is

Knowledge and Modality

In order to avoid the attack Gettier launched against the classical tripartite definition that it seems reasonable to focus attention on modalised theories of knowledge. Whatever other requirements a theory of knowledge must meet, it has to account for the fact, hinted at above, that an accidentally true belief does not amount to knowledge. In order to know, one has to get things right in a manner that secures truth, not merely in the actual situation, but also in a range of hypothetical situations. For example, for Foucault to know what time it is, we would demand of Foucault that he got things right, even had things gone slightly differently and his clock had stopped 5 minutes later than it in fact did.

No doubt the most influential modalised theory of knowledge has been the so-called 'tracking account' developed by Dretske (1970) and (1971) and Nozick (1981). The basic idea behind this theory is that mere justification is not appropriate for gaining knowledge. The supporting justification has to be of a certain kind. In particular, it must be constituted in such a manner that were it to be constituted in an identical, or very similar manner in slightly modified scenarios, it would still hit the target, as it were. When we direct our beliefs towards worldly truths justification serves as a guide. The metaphor of belief as an arrow is helpful here. Every time an arrow is discharged, its direction and destination are determined by a complex and interacting system of steering devices; the initial angle of the arrow relative to the intended goal, the tightness of the bowstring, the design of the steering feathers, the weight of the arrow, etc. And whether the arrow hits the intended target depends, to a large extent, on the constitution of this steering device. But not exclusively, for there are also certain device-external factors involved in determining the arrow's direction, such as the force of the wind, the mobility and size of the target, etc. Now, in order to be a master of archery it is not sufficient merely to be capable of manipulating the steering device such that the arrow hits the target under a particular specifiable external condition. It does not even suffice to be capable of manipulating the steering device in such a manner that the arrow hits the target repeatedly as long as these specified circumstances are held constant. Even a blind person might accomplish that given he has been appropriately trained. However, the slightest change in device-external surroundings would result in such a blind person missing the target, no matter how otherwise skilled he is.

A true master, therefore, is one who, in manipulating the steering device, takes such potential changes of external factors into account. His manipulation of the steering device is so constituted that, were it constituted identically, or very similarly, under varying circumstances, he would still hit the target. The manipulation of the steering device must be *sensitive* to the device-external factors. Likewise, the thought is, in modal theories of knowledge, that the processes by which a belief is formed have to be so constituted that, were they to be identically,

by no means meant to imply that this issue between epistemic internalism and externalism is foreclosed.

or very similar constituted, in slightly modified hypothetical scenarios, they would still lead to a true belief. A justification for the claim to know the time consisting in consulting a stopped clock does not satisfy this requirement. Granted, the evidence so obtained might lead to a true belief. It might even do so repeatedly given it always is constituted under these particular circumstances (that consultation occurs at a particular time of the day). But were Foucault to base his belief on this evidence in slightly different circumstances, for example circumstances in which the clock stopped 5 minutes earlier than it did, his belief would no longer be directed towards the truth.

This way of presenting matters raises the following question: what is it for belief-forming procedures to be constituted in a certain way and, *a fortiori*, what it is for justification to be identically, or very similarly constituted in various hypothetical situations? And it is in attempting to answer this question that Dretske's and Nozick's theories ultimately reveal their deficiency. Since, without resources to answer this question a modal epistemology loses its explanatory power, we take it to be one of the major tasks of our investigation to search for a suitable response.

One of the issues that this question raises concerns internal vs. external standards for justification. Someone in favour of internalism will argue that every relevant feature of the constitution of an agent's justification must be accessible to that agent's cognitive apparatus—must be, as it were, open to view for the agent. In the clock case an internalist would thus claim that the features relevant to a discussion of what constitutes Foucault's justification are such features, as it appear to Foucault that there is an apparently well-functioning clock in front of him and, furthermore, that the perceptual conditions seem, to Foucault, to be well suited to investigating what time it is by consulting the clock. An externalist, on the other hand, would claim that these internal constituents to Foucault's justification, although relevant, may be augmented by factors exceeding the scope of what is immediately open to view for Foucault. An externalist would thus insist that Foucault's justification must be constituted also by the facts that there really *is* a clock in front of him, that it *is* well functioning (if it is) and that perceptual conditions for consulting the clock really are appropriate—all this independently of how matters might or might not appear to Foucault.

Here again the analogy between belief and an arrow is helpful. I said that the archer's manipulation of the steering device has to take into account potential changes of device-external circumstances. The manipulation must be constituted in such a way that it is sensitive to these device-external factors. The archer must somehow, consciously or unconsciously, read these device-external factors—wind, speed of the target, etc.—and calculate how the steering device should be calibrated accordingly. Competent archery can be considered a two-way system. It involves both a passive intake of external clues and a controlled (outgoing) dispatch of the arrow. The archer receives clues about factors external to the steering device, manipulates the steering device accordingly, and becomes, thereby, capable of sending the arrow towards its target. The archer can be said to hold a repertoire of default adjustments suited for various standard situations—

default adjustments which he is continually ready to make to meet relevant external factors.

Is there any analogous structure to be found in the case of belief formation? Well, in order to answer that question it is important to clarify what the analogous components are. I suggested the arrow may be compared with belief and the steering device with justification. But what is the counterpart of the factors external to the steering-device? What is the counterpart of the circumstantial variables, modification of which influences whether a particular adjustment of the steering mechanism guides the arrow to the target? In order to answer these questions we need to determine what the circumstantial variables are, the modification of which influences whether a given piece of justification issues in a true belief. And those that naturally come to mind are the agent's epistemological settings. If we accept the internalist account of justification we can say that a piece of evidence might guide an agent to the truth in one setting, and maybe even say that it might always guide the agent towards truth in that setting. But we can still claim that, were this setting to be slightly modified, that piece of evidence would no longer be trustworthy. For example Foucault, since he has recently consulted his clock, is (internally) justified in believing what time it is. And if he is put in a setting which has it that clocks are well functioning, that piece of (internalistic) justification will invariably (or almost invariably) lead him to the truth. But if he were to be put in a different setting which has it that clocks are not generally reliable—were he for instance in a museum for old clocks—then that particular piece of evidence might not be any good for him in sorting out what the time is. And if Foucault misreads his own settings, if he is not sensitive to clues about which settings he happens to be in, the evidence might even become manifestly misleading.³

Epistemic Settings

But, you might object, surely Foucault cannot be mistaken about his own settings. It is platitudinous that he is aware of what settings he happens to be in. Of course he knows whether he happens to be at home perceiving his own clock or in a

³ It is no easy matter to state, in any formal manner, what settings are. The core idea underlying the notion is that almost invariably there is a gulf between a given piece of (internalistic) evidence (as for example when it appears to Foucault, from consulting his clock, that it is 6 o'clock) and the state of affairs supported by this piece of evidence (for example it is 6 o'clock). We could, if we wished to, state a very long list of collateral states of affairs that have to obtain in order for that piece of evidence to be veridical. (For example Foucault is not sleeping, conditions are suitable for perceptual investigation, the clock is well functioning) If the entire list holds true, circumstances are conducive to the investigation in question. This fictive list would state the epistemic settings relevant for the particular investigation. 'Settings' is thus an external notion; the world's contribution to a successful epistemic investigation, as it were. However, a cognitive agent who forms a belief based on a particular bit of evidence will, as a matter of course, believe the world to be co-operative in this sense, consciously or subconsciously.

museum for old clocks. Likewise he would be extraordinarily inattentive if he hadn't noticed whether his clock is at least moderately reliable. And more generally it seems that his beliefs regarding his other settings—that he is not asleep, etc.—‘second-order’ beliefs as we might phrase them, will always be appropriate, and thus he always will be capable of calibrating his (internalistic) justification relative to the right settings.

This is at least how we usually conceive of these matters. As in archery, it makes sense to consider our epistemic relation to the world as a two-way system. We are continuously confronted with all sorts of raw uncensored evidence, ‘first-order’ evidence, which provides for an, at least initial, reason to hold a particular belief. But in addition we possess second-order beliefs in virtue of which we access and rank our first-order evidence. We take ourselves to have some correct opinions regarding the reliability of various pieces of first-order evidence. Indeed, the world would be an awful place to be if we didn't. This is very clearly illustrated in Douglas Adams' science fiction novel *Dirk Gently's Holistic Detective Agency*. Adams sets his story in a future society in which robots have taken over all the dull and laborious functions in life. Robots do all the monotonous labour at the factories. They do the dishwashing, the car washing, etc. Robots have even taken over the television watching (the most advanced of the tele-robots being capable of watching as much as fourteen channels simultaneously). The most recent, state of the art development, is a brand new robot, which has been devised to care for religious duties. These ‘electronic monks’ believe for you! And they believe everything they see or hear without any sense of critical reflection. One day, however, things go terribly wrong: by mistake one of the electronic monks is swapped with a tele-watcher robot. After just 5 minutes the poor electronic monk suffers a severe nervous breakdown: believing fourteen distinct television channels is more than even a 160 MB electronic monk can capacitate (the 160 MB being calculated as sufficient capacity to cope with five (inconsistent) major religions, sixteen complete party political broadcasts and Ronald Reagan's autobiography). Adams' scenario trades on the fact that evidence is despairingly rarely unequivocal. Nearly every time we possess a warrant for believing some belief P we also possess some contradictory warrant for believing $\neg P$. For example, while testimony from some informant warrants us believing she is in a particular state of mind—say cheerful—close observation of her behaviour warrants us believing she is not. One of life's major challenges is to assess the reliability of such contradictory warrants and to form beliefs accordingly.

Electronic monks are incapable of doing this. We, on the other hand, although not infallible in that respect, have some degree of competence. We are capable of deciding between contradictory warrants and, generally, of accessing all sorts of first-order warrants and ranking them mutually with respect to credibility. That is, our cognitive functioning is not restricted to merely gathering all sorts of evidence and more or less blindly forming beliefs accordingly. We are, in addition, capable of critically assessing such justification. Some evidence is considered amenable to all sorts of distortion. For example, when a television advertisement tells us a particular brand of washing powder by far exceeds its rivals in efficiency and

gentleness we automatically make certain reservations regarding this information. We envisage several possibilities for error: circumstances under which the (default) evidence is non-veridical. Other sorts of evidence, on the other hand, are considered highly plausible. For example, were we to discover, by inspection, that the car isn't parked where we left it 2 hours ago we would tend to assess this piece of observational evidence as highly reliable. In this case we would, on most occasions, not even bother to consider possibilities of error: scenarios in which the evidence for some reason might be misleading.

At some universities you will take it as evident that one of your fellow students had borrowed a car if you see her driving one. At other universities, however, seeing a fellow student driving a car is pretty good evidence that she owns one. And there are, presumably, universities at which the same evidence can be safely taken to indicate that a fellow student seen driving a car recently has stolen it. Settings just vary. And to get them right, and hence adjust our first-order evidence appropriately, is crucial to competent cognitive functioning.

Second-order beliefs can perform the function of adjusting first-order evidence appropriately. And second-order evidence secures that the second-order beliefs do not do so blindly. But what is second-order evidence? It can be understood as a grasp of one's own settings, i.e. an ability to determine whether the present setting is one in which a given piece of first-order default evidence is to be considered reliable. In archery an apparently favourable trimming of the steering device might send the arrow a few inches above the target. If it does so repeatedly the skilled archer will learn a lesson from this: some factors external to the steering device, a strong head wind say, influences the arrow's curve towards the target. Hence, strong head wind must be accounted for by aiming a bit lower than usual. Likewise some apparently favourable observational evidence might suggest some person has a particular character—is magnanimous and honest, say—but beliefs to just that effect might, disappointingly, prove fallacious. If it happens repeatedly that such first-order evidence turns out to be fallacious in a particular social environment the skilled cogniser will learn the lesson and aim a bit lower when forming beliefs regarding other people's characters based on observable evidence. The skilled cogniser will thus learn to read his own settings more carefully in the future: when situated in a particular social environment, possibilities of error will be considered likely in the light of, otherwise reliable, evidence. Ultimately second-order evidence—competence to adjust one's first-order evidence to the appropriate settings—is thus based on inductive inferences from past experiences.

How radical a modification must such second-order evidence be capable of accounting for? Well, in archery we are inclined to say that every practitioner, even a Zen master, has his limitations. In a situation with gale force winds and in which the target is very small and moving at very high speed it is highly probable that even a Zen master would miss his target. And if he is so exceptionally skilled that he would hit even under those circumstances we could readily add some further potential modification, such as momentary suspension of the gravitational forces, the target turning invisible, etc., under which he certainly no longer would hit the target. But that, in itself, does not challenge his being an archery master. To be an

archery master his manipulation of the steering devices only has to be so constituted that he will hit the target within a range of *reasonable* hypothetically modified situations. We demand of him that his manipulation of the steering device is sensitive to a reasonable range of device-external factors such as, for example, the general wind conditions. But we do not expect of the archer that he should be capable of taking into consideration such far-fetched hypothetical changes in device-external conditions as for example, the target being removed at the last moment by a guardian angel.

Similarly, in the case of belief formation, we must admit that potential modifications of circumstances are conceivable as a result of which even the most promising justification will be misleading—even for a skilled cogniser who has got things right regarding his own settings. The clock case illustrates this point. Foucault's justification for his belief is, *prima facie*, strong. He has just consulted his clock. In addition, Foucault has a qualified second-order belief regarding this sort of evidence. He is aware of his own settings. He knows his clock is reliable and hence that consulting it provides a strong warrant for believing what time of day it is. He is like the archer who has read the wind-conditions correctly, and manipulated the steering device accordingly, but whose arrow has been led astray by a sudden change in the wind (cf.: *this* time the clock isn't working properly), but, miraculously, which has been put back on the right course by another, likewise unpredictable, windfall (cf.: Foucault consulted the clock at *that* particular time) and thus eventually hits its target. I said Foucault is lucky. And he is. But he is, at the same time, extraordinarily unlucky. Although he has a pretty good grasp of his own settings, and has a suitable second-order belief regarding the reliability of consulting-the-clock evidence, he is wrong about his own settings in this particular case. He has the second-order belief that his clock is working properly. But it is not.

First-order evidence of the sort available to Foucault, namely looking at the clock, does not amount to knowledge under conditions in which the relevant second-order beliefs are mistaken, that is the conditions described in, or similar to those in, the clock case. This is rather trivial. However, the reciprocal claim, that whenever the relevant second-order beliefs are veridical, any (first-order) evidence counts as knowledge, might shed some light on our initial question: what it is for a warrant to be constituted in the same, or very similar manner, under slightly modified circumstances. An initial, although by no means exhaustive approach to this question might be the suggestion that a warrant for P is identically, or very similarly constituted, if it is acquired (i) in a context where the relevant second-order beliefs P' are held constant; and (ii) in a context where the truth-values of those P' are held fixed. According to this suggestion, had Foucault consulted his clock a couple of days earlier while it was still functioning properly, his evidence would not have been constituted the same way as in the original case: although (i) above would be satisfied, (ii) would not. (i) would be satisfied since Foucault would still have believed that his clock was working properly if he had consulted it a couple of days earlier. But (ii) would not be satisfied since the relevant second-order belief—the belief that the clock is working properly—actually would have been veridical if Foucault had consulted the clock a couple of days earlier.

Likewise, if Foucault is to look at a clock in a museum for old clocks, his evidence is again not constituted in the same way as in the original case: this time because (i) is not satisfied, although (ii) is. This time (i) is not satisfied since Foucault would not assume that an old clock in a clock museum is working properly. But (ii) would be satisfied since a second-order belief to just that effect, as in the original case, would have been fallacious.

Fallibility

I said common sense has it that we always have opinions regarding our own settings, that we possess second-order beliefs, and, furthermore, that it is close to platitudinous that these second-order beliefs, by and large, are veridical. That, of course, is not to claim that all our second-order beliefs are infallible. Some of them, such as opinions on whether we happen to be at home or in a museum for old clocks, are presumably infallible according to common sense. Others, such as the belief for example that our clock is working properly, might be assigned merely a high degree of probability. Although it is a matter of course that we know whether a familiar clock is reliable, it is still in conformity with common sense that error occasionally sneaks in. However, in either case the relevant second-order belief is detectable, as it were. Indeed it is essential to the idea of a second-order belief containing such modal notions as 'reliable' that they must occasionally be fallacious, if only very rarely. And it is natural to suppose that it is in virtue of the fallibility of our second-order beliefs, and the resultant rare pitfalls, that talk about being right and wrong makes sense in the first place. Far from being an obstacle, occasional error regarding one's setting is, in this sense, rather conducive for knowledge acquisition. Justification need in any case not take account of *this* kind of setting variation in order to count as knowledge. It suffices that the settings have been properly read with a certain degree of probability (whatever exactly that degree may be). It is, in other words, not necessary to rule out every possibility of an error scenario.

However, another sort of circumstances is also conceivable which renders first-order evidence misleading. To your big surprise you discover your car is not parked where you left it a few hours earlier. Your second-order belief here is that you are not suffering any sort of mental or perceptual disorder, that the viewing conditions are normal, that you are capable of recognising your car if confronted with it etc. Each one of these second-order beliefs might be fallacious: your 'friends' might have played a trick on you and, in your absence, painted your car a different colour, so that you now don't recognise it as yours, or they might have covered it behind a cleverly painted screen, or they might presently be hypnotising you into merely believing your car has been removed. Each of these three error-scenarios differs from the clock case in that you don't consider yourself to have to rule them out 'with a certain degree of probability'. You don't consider yourself to have to rule these far-fetched possibilities out for the simple reason that, under normal circumstances, they don't need to be taken into consideration at all. Under