Calibrating Introspection

Spener, Maja

DOI:

[10.1111/phs.12062](https://doi.org/10.1111/phs.12062)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Spener, M 2015, 'Calibrating Introspection', *Philosophical Issues*, vol. 25, no. 1, pp. 300-321.
<https://doi.org/10.1111/phs.12062>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is the peer reviewed version of the following article: Spener, M. (2015), Calibrating Introspection. *Philosophical Issues*, 25: 300–321. , which has been published in final form at <http://dx.doi.org/10.1111/phs.12062>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

Checked Feb 2016

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Calibrating Introspection

Maja Spener (University of Birmingham)

Published in *Philosophical Issues*, Vol. 25, Issue 1, October 2015. This version of the paper is only a **draft**. For quotation, please consult the published version.

1 Sceptical Worries

The epistemic credentials of introspection are often overestimated by philosophers and psychologists in their theorizing. This point has been argued in recent literature (e.g. Dennett 1991; Schwitzgebel 2008, 2011; Spener 2011a, 2011b). Yet a sceptical position on introspection which recommends its wholesale rejection as a source of knowledge is not credible (see Bayne & Spener 2010). We can and do know things about our own minds by introspection. A sensible view on introspection is therefore moderate in its scepticism. It is surprisingly hard, however, to make a principled case for such a position. In this paper, I put forward and defend a moderate scepticism about introspection. I show how this position yields the resources to pin down the scope and content of introspective knowledge. A major attraction of the view is that it provides a strategy for solving a problem which has vexed the scientific study of consciousness from its inception.

The conception of introspection at work in this paper is deliberately liberal and somewhat loose. According to it, introspection is a distinctively first-personal way of forming judgements—a way of forming judgments ‘from the inside’ as it is sometimes put. Introspection contrasts with other ways of forming judgements, such as forming judgements by apriori reasoning and forming judgements by endorsing or rejecting empirically acquired information. There are no particular restrictions on what a first-personal way of forming judgements involves, e.g. whether such judgement formation involves explicit or deliberative reflection of some sort, or whether the resulting judgments are conscious, or even whether there is more than one type. There are also no restrictions on the targets of judgements formed in this way.

This liberal conception of introspection will end up counting many, and very likely too many, judgements as introspective. Moreover, given that our views of apriori reasoning and other judgement-forming ways are typically imprecise themselves, a conception of introspection drawn solely in terms of a contrast with them will inherit this lack sharpness. Ultimately, an adequate view of introspection should not suffer from these shortcomings and so will be less liberal by positing significant restrictions of the sort mentioned above in addition to

contrasting it with other kinds of judgement formation. Different restrictions lead to very different conceptions of introspection. In this paper, though, I want to keep everyone on board by staying relatively neutral on how to flesh out one's view of introspection beyond the central contrast with other ways of forming judgement. The general view I put forward about moderation in one's scepticism is independent of particular conceptions of introspection, but the latter do influence what exactly one thinks the content and scope of introspective knowledge is. I will explain this along the way.

Having said that the operative conception of introspection is liberal, let me nonetheless introduce some fairly standard examples. Paradigmatically, introspection includes first-person reflection on one's conscious mind. I might judge in this manner that I have a pounding headache, for instance, or that I feel ravenously hungry, or that it looks as if there is something green over there. All these are introspective judgements about current conscious mental states: having a terrible headache, feeling ravenously hungry, and having a visual experience as of something green over there.

Ordinarily, when I introspectively judge that I have a pounding headache or that I am feeling a pang of guilt, say, I take myself to know that I am in those mental states. We normally take ourselves to know about our minds by introspecting. However, we also make room for the idea that there are introspective judgements that do not, for various reasons, figure in states of knowledge. If one reflects too hastily, for instance, one might mistakenly think that one is feeling exhausted and low when really one is feeling very hungry. Or, if one is under the influence of drugs that impair cognitive faculties, one's introspective judgements cannot figure in states of knowledge even if they happen to be correct. Intuitively, then, there are introspective judgements that figure in states of knowledge—let's call them 'good' introspective judgements—and there are those that don't—let's call them 'bad' introspective judgements. This terminology is meant to mark nothing other than that these judgements figure in, or fail to figure in, knowledge.

Scepticism about a given range of introspective judgements is the view that none of the judgements in that range figure in states of knowledge, that none of them is a good introspective judgement. The most extreme version, global scepticism about introspection, says that there are no good introspective judgements at all, that we do not have any knowledge via our first-personal ways of forming judgements. Global scepticism about introspection constitutes a radical departure from common sense. The intuition that in the ordinary course of events we know about having headaches and feelings of hunger on the basis of introspection is strong and widespread. It is a sign of just how ingrained this common sense assumption is that even the fiercest proponents of scepticism about introspection usually cannot stomach its extreme version. At some point even they tend to give way to the intuitive pressure that introspection affords some minimal knowledge of very general facts about our own minds (e.g. Dennett 1991: 68–9; Gertler 2010: 118).

In light of this, any plausible scepticism about introspection must be moderate, entailing that some introspective judgements are good while others are bad. This much is fairly uncontroversial. Acceptance of moderation, though, tends to rest simply on the deep-seated common sense intuition that introspection delivers some knowledge and the thought that giving it up is simply too costly. But such intuition-driven moderate scepticism is unsatisfying. There are two main worries: firstly, it is unclear how to weigh the force of reason-driven scepticism against that of intuition-driven moderation. This renders the view unstable because it is constituted by sceptical reasons pushing towards global scepticism and

intuition pushing away from it. Secondly, following on from the first, intuition-driven moderate scepticism does not yield a principled way to locate the scope and content of introspective knowledge. It just says that there are some good introspective judgements, but does not yield any more specific verdicts on what sorts of things we can know introspectively. Thus, while the general moderate position will be readily agreed to, controversy will immediately flare up when drawing the line between good introspective judgements and bad ones.¹ More seriously, such a view will not offer a way to tell when introspection yields trustworthy data.

In this paper I provide a reason for moderation in our scepticism about introspection, beyond any intuitive support, deep-rooted though that may be. In consideration of the worries just articulated, this reason seeks to satisfy two desiderata. First, it aims to anchor common sense belief that we have introspective knowledge by revealing the source of our intuition which so powerfully motivates moderation in scepticism about introspection. Such anchoring would stabilize the overall position because it would allow appropriate weighing-up of negative and positive forces concerning introspective knowledge, all the while not leaving our theoretical support for moderate scepticism out of whack with our intuitive support for it. Secondly, the reason for moderation seeks to offer a strategy for locating the scope and content of introspective knowledge as such.

This project is different from well-known philosophical concerns with the nature of self-knowledge. The latter deal with what introspective knowledge is and how, if at all, it differs from other kinds of knowledge. For example, introspective knowledge is variously claimed to be observational in some sense, or to be constitutive of rational first-order belief, or to involve following certain basic epistemic rules (see Gertler 2010, Shoemaker 1994 and Byrne 2011 for representative samples of such views). I am not concerned with such issues. My aim is to provide a principled reason—beyond common sense conviction—to believe that we have some introspective knowledge, and to do so in the face of serious sceptical worries. The view I develop tells us when it is rational to decide that certain introspective judgements are good. It secures a rational basis for our widespread and ordinary belief in introspective knowledge by showing how that belief is grounded in a certain kind of explanatory practice. In addition, this grounding anchors common sense conviction about such knowledge and thus vindicates and explains common sense intuition about this matter.

Moreover—and this relates to the second desideratum mentioned above—once the rational basis for introspective knowledge is understood, it yields a strategy for settling on its proper scope and content. The main incentive behind this second desideratum is methodological, concerned with the evidential status of introspective judgements within philosophical and scientific theorizing about the mind. The strategy I offer promises a way of solving what is sometimes called ‘the problem of calibration’ for introspection:

A crucial problem for the theory of introspection is to fix its range of epistemic success. This is the problem of calibration, which arises for any scientific instrument and cognitive capacity. I would subdivide the problem into two parts. One would seek to specify the operational conditions under which introspection is sufficiently epistemically successful. The second would seek to specify the propositional contents for which it is epistemically successful, i.e. the types of mental

¹ Elsewhere I have engaged in such controversy: I have argued, for instance, that we should be highly sceptical of certain introspective judgements routinely considered good by many philosophers (Spener 2011a; ms). But these sorts of cases for scepticism about introspection are piecemeal, pertaining to particular disputes about specific ranges of introspective judgement.

or cognitive descriptions for which introspection is prone to be accurate (assuming operational conditions are right). (Goldman, 2004: 14)²

As Goldman points out, every source of knowledge is subject to a calibration question about its range of epistemic success. But introspection is usually held to be particularly problematic in this respect (e.g. Gertler 2011). The main reason for this is that the deliverance of introspection is first-personal and seemingly private, and it is hard to see how it can be checked effectively against information from other, non-introspective sources (e.g. Dennett 1991: 55). At the same time, there is mounting evidence that introspection is epistemically suspect in that it generates mistaken judgements as well as radical disagreement among introspecting subjects. Calibration, then, seems badly needed but beyond our reach. Having said that, by most people's lights introspection is the primary source of knowledge about our conscious experience. Hence, it is the primary source of first-person data which feeds a scientific investigation of consciousness. But if—as sceptics typically argue—introspection is an inherently and pervasively epistemically unsuccessful source of data, no method relying upon such data can be scientifically respectable. In that way the problem of calibration for introspection threatens to put a genuine science of consciousness beyond our reach as well. If, then, there is any chance for optimism about the future of a scientific study of consciousness, in the face of extensive sceptical worry about introspection, it must lie with a solution to the calibration problem.³ The hope is that we can know enough by introspection to provide us with a sufficiently firm finger-hold on the conscious phenomena we wish to study.

In putting forward a strategy to pin down the scope and content of introspective knowledge, I make a case for optimism about the prospects for a scientific study of consciousness. In outline, I propose that introspective judgements can be rationally held to be good—i.e. rationally held to figure in states of introspection-based knowledge—when doing so helps explain our having introspection-reliant abilities. In the next section I illustrate the overall shape of this proposal by developing an analogous one for a distinct source of cognitive contact with the world, visual perception. I do so because we are generally more familiar with the visual case and it will be helpful to tap into that familiarity when setting out my view about introspection. In section 3 I turn to introspection, first setting out examples of relevant abilities to be discussed, and then using them to apply my proposal to introspective judgements. In section 4 I show how this proposal anchors common sense intuition about introspective knowledge. I also explain how the verdicts rendered by the proposal (and, in the end, common sense intuition about introspective knowledge) are sensitive to correction and revision in light of relevant empirical results. In the last section I explain how my proposal yields a strategy to solve the calibration problem for introspection.

2 An Analogy with Visual Experience

Ordinary cases of seeing the world involve conscious mental states. These states occur when something visually appears, or looks, to be certain ways. When I cast my eye around the room, for example, it looks to me as if there is a grey chair before me, and something orange and round in the corner. These conscious states are visual experiences; for me to be in one of them is for me to be visually presented with something. Moreover, visual experiences

² I am substituting the phrase 'epistemically successful' for 'reliable' (*mutatis mutandis* for their cognates) because my discussion is framed in terms of knowledge. Goldman (and others) cast these issues in terms of reliability instead. Since the notion of reliability is not used consistently across philosophers or psychologists, and within epistemology it is subject to a host of problems, I prefer the more everyday knowledge talk.

³ Granting that introspection is a necessary source of data for investigating consciousness.

intuitively come in two varieties: those that figure in states of genuinely seeing the world aright—‘good’ visual experiences—and those that do not figure in states of genuinely seeing the world aright—‘bad’ visual experiences. Among the latter are those that figure in various kinds of hallucination and illusion.⁴

Is there a principled way to judge, on a given occasion, whether a certain range of visual experiences is good? The starting point for showing that there is, is the role that we ordinarily take visual experiences to play in our agential lives. More specifically, my proposal takes advantage of an explanatory link which is very much part of common sense practice between having good visual experiences and having ordinary abilities.

Consider the ability to find one's way home. With a normally-sighted person, we find it natural to explain the exercise of this ability by appeal to her seeing her environment correctly, e.g. her seeing the corner where she must turn left, where to cross the street, etc. This explanation of the exercise of an ability appeals to visual experiences assumed to be good—i.e. those figuring in states of genuinely seeing the world aright (at least in relevant respects). It is the fact that they are good that makes them explanatorily efficacious in the first place. Bad visual experiences would not contribute to an explanation of the exercise of the ability to find one's way home.⁵ Of course, one might hallucinate in a way which, by sheer luck, proved helpful in getting home, but such a case would not constitute exercise of an ability. It would be a matter of luck, a fluke accomplishment. When good visual experiences are needed for the explanation of an event such as finding our way home our explanatory target is not a one-off event, it is the exercise of an ability. So, to count as an exercise of the ability to A, a given A-ing must satisfy some anti-luck conditions: the performance should be robust enough to be repeatable in sufficiently similar situations. We are after an explanation of the manifestation of an action-related property of the agent—having it in one's power to find the way home from a certain place, for example.

A complete explanation of this sort would not merely appeal to good visual experiences. There are other important abilities and kinds of states of the agent which play a role, such as being able to walk and to remember things normally. The fact that there is a certain organized interaction among exercises of these abilities also importantly contributes to the explanation of how an agent exercises the target ability. Nonetheless, my point here turns on the role good visual experience plays: with normally-sighted people, such experience makes a significant contribution to the overall explanation of their exercise of the ability to find their way home.

A similar story applies to the exercise of other visually-based abilities: running through a tightly-planted forest, climbing a certain mountain, playing football, etc. In common sense practice an explanation of the exercise of these abilities by a normally-sighted person will make crucial appeal to visual experiences figuring in states of seeing the world aright. Such an explanation will make appeal not only to visual experience as such, but to states of seeing the world aright (in certain respects). Good visual experiences are a key ingredient in these sorts of explanations of the typical manifestation of such abilities. This is because they

⁴ I am using the phrase ‘figuring in’ to leave it open whether good visual experiences simply amount to states of seeing the world aright or whether they are merely one ingredient in the overall state of seeing the world aright.

⁵ This must be qualified a little. Bad visual experiences can play a role in such explanations if they are near enough good. For example, if I have a slight astigmatism producing visual experiences with a mild perspectival distortion, the latter could still serve in explanations of exercises of my ability to find my way home. The bad aspects of such visual experiences are irrelevant to that explanation, the experiences can count as good, or good enough, in these contexts.

provide an information-based, safe connection to the environment that allows one to coordinate actions into successful navigation of obstacles. As I indicated above, explanations of this type are natural in that we regularly use them to make sense of what we can and cannot do in particular situations. This explanatory practice is by and large very successful.

In light of this explanatory connection between good visual experiences and ordinary abilities, here is a first pass at a proposal about rationally holding visual experiences to be good. We can rationally hold a given range of visual experiences to be good when treating them as such helps explain the exercise of a visual-based ability.

While on the right track, this proposal is too restrictive. There are plenty of cases in which someone is not engaged in any relevant exercise of an ability, but where it is rationally appropriate to attribute good visual experiences nonetheless. For example, someone might be trying to exercise a relevant ability and not be successful on the occasion, where her failure is not due to a problem with her visual experiences. Perhaps the road has been dug up and she cannot get across the street, thereby thwarting her attempt to get home. A proposal that fails to recommend rationally holding visual experiences to be good in such situations will turn out to be silent about a vast amount of cases where it is intuitively plausible that we have good visual experiences. This would make it unsuitable for providing a rational basis for the common sense belief that routinely we have such visual contact with the world, as it would underwrite that belief in only a small subset of cases. The lesson to be drawn here is that a suitably inclusive proposal must not tie attribution of good visual experiences too closely to their explanatory role in the *exercise* of relevant abilities.

Consider Lazy Person. She lounges in a garden chair, looking at her cat sleeping on the bench next to her. Even though Lazy Person is not doing anything much at the moment, we are happy to say that there are many things that she can do. For example, she is able to sing, to drive, to speak a certain language, to let her cat into the house, to shoo her cat away, to give her cat food. These are all abilities that Lazy Person has but is not exercising at the moment. A second pass at a proposal for rationally holding visual experiences to be good shifts the target of explanation away from exercise of an ability, focusing instead on the *possession* of an ability. The core thought is that holding a given visual experience to be good helps explain the having (and not merely the exercise) of certain abilities. For instance, holding Lazy Person's current cat-experience to be good helps explain her having certain abilities which involve the white cat in front of her. Lazy Person's seeing the cat helps to put her in a position to shoo the cat away, to give it food, to let it into the house, and so on, in the situation she is in. By appeal to her good cat-experiences, we can explain Lazy Person's possession of a range of abilities indexed to her current environment.

This is a familiar explanatory practice. Why is Lazy Person able to shoo the cat away right now in her current situation? In part because she sees it. Note that this is not a causal explanation: seeing the cat does not cause the ability to shoo the cat away, or cause the exercise of that ability. Rather, the relationship is constitutive, i.e. seeing the cat is part of what makes it the case that Lazy Person has the ability in question. More generally, we can and routinely do explain a subject's possession of such environment-indexed abilities—of what she has it in her power to do in her situation—by attributing to her some cognitive access to her environment. For these kinds of abilities, such cognitive access is normally provided via good visual experience. In turn, this explanatory role of good visual experience yields a principled way to judge, on a given occasion, whether a certain range of visual

experiences is good: we can rationally regard visual experiences as good when doing so helps to explain the possession of environment-indexed visual-based abilities.⁶

3 Holding Introspective Judgements to be Good

Let us return to introspective judgements. In close analogy with our discussion of visual experience, I suggest that we can find a principled reason to hold introspective judgements to be good when we consider the explanatory role of introspective judgements with respect to our abilities. Just as before, for the proposal to be sufficiently inclusive, the explanation must target the possession rather than merely the exercise of relevant abilities. Before discussing the proposal itself, however, I want to introduce the kind of abilities which figure in these explanations. This is best done by spelling out what is involved in their exercise—so although the proposal I put forward concerns possession of such abilities, I shall first talk about their exercise.

Consider the following three cases, in each of which it is plausible that we can explain a subject's exercise of an ability by appeal to good introspective judgement.⁷

Focusing Binoculars: You are outside bird watching. You point your binoculars at a bird in the distance, but you cannot see anything clearly. Being an avid birdwatcher you immediately fix that: you focus your binoculars. You peer through the lenses and make judgements about whether or not things look fuzzy, adjusting the lenses until things look sharp; then you stop adjusting the lenses. A natural explanation of your exercise of the ability to focus your binoculars appeals to judgements about how things look. These judgements concern some of your mental states, namely some of your visual experiences and you arrive at them by introspecting how things look. You exercise an ability to focus your binoculars by using good introspective judgements about whether or not things look fuzzy.

Ordering Pizza: You are ordering a pizza and you must choose between a large and a small one. While you don't want to spend more money than you have to, you also don't want to be hungry at the end of the evening. Having gone to restaurants before you make the decision successfully. An important part of doing so is that you gauge your hunger. You judge how hungry you feel and that plays a key role in deciding between a large or a small pizza. In this way, a natural explanation of your exercise of the ability to order the right-size pizza appeals to judgements about how hungry you feel. These judgements concern some of your mental

⁶ I said that this sort of explanation is part of ordinary practice. It also fits the model of a more general approach to psychological explanation, i.e. explanation of a psychological property by functional analysis. One important goal of psychological explanation is to explain possession of psychological capacities. A functional analysis does so by explaining a given psychological capacity in terms of constituent capacities possessed by components of the overall system, plus a specification of their programmed exercise. The explanation thus involves showing how the exercise of one sort of capacity depends upon the systematic joint exercise of other sort of capacities (see Cummins 1975, 1983, 2000).

⁷ Two caveats upfront, though. Firstly, these cases are meant to represent a manageable list of compelling examples. I do not mean to suggest that the list exhaustively illustrates the kinds of cases I have in mind. Common sense underwrites a large, heterogeneous class of abilities of this kind. Secondly, recall that the conception of introspection at work in this paper is ultra-liberal. According to that conception, introspection is a distinctively first-personal way of forming judgements, contrasted with other ways of forming judgments. This conception ensures a wide array of possible examples, not all of which will seem compelling to everyone. Much depends on which further restrictions one's ultimate conception of introspection imposes on first-person judgement formation. Different examples thus work for differently filled-out conceptions of introspection. As I said in the beginning, the proposal I am developing is independent of the specifics of particular examples. The following cases are the ones I am using to set out my view, but I would happily substitute others.

states, namely your feelings of hunger or appetite—and you arrive at those judgements by introspecting how hungry you feel. That is, you exercise an ability to order the right-size pizza by using good introspective judgements about whether or not you are very hungry.

Slowing Down: You are in your office at 4.45pm after a very busy day. You have to take some papers over to the Dean's office, pick up your daughter from day care, drop off some signed documents at the bank and get your car from the workshop all of which close at 5.45pm, and do some shopping for dinner. Getting all of this done is important yet does not immediately seem feasible. You are feeling incredibly stressed and try to figure out how to get everything done while sprinting out of the office, rapidly thinking things like 'well, if I drop off the papers first, and then take the bus to the workshop, I'd just about make it to day care by 5.45pm—oh, but the shopping, I'd have to drive all the way back to the shop and we'd only be home very late, so maybe I'll go to get my daughter first and take the bus with her, no, that won't work—oh, and what about the papers...' In the end, you realize that you cannot carefully and reliably plan in this frame of mind. You stop and make yourself think more carefully and systematically about how to get everything done most efficiently in the time you have. A good explanation of your exercise of this ability to slow down your thinking, to ensure that you are planning scrupulously in this situation, appeals to judgements that you are thinking things through properly. Such judgements are part of your effort to reason well given that the stakes are high. Among other things, you judge that you are thinking about the options carefully—i.e. slowly and systematically enough to avoid making mistakes—and you arrive at such a judgement by introspecting the pace of your train of thought.

All three cases feature ordinary explanations of ways in which a given ability is exercised by a normal subject. According to these explanations, the abilities have in common that their exercise involves judgements that are got by introspection. I call such abilities 'introspection-reliant' in that their exercise can (in part) be constituted by good, i.e. knowledgeable introspective judgements. The ability to focus one's binoculars, the ability to order the right-size pizza, and the ability to slow down one's thinking for careful planning are all examples of introspection-reliant abilities. I must emphasize, though, that in classifying abilities as introspection-reliant, I do not claim that their exercises are *always* best explained by appeal to introspective judgements. Exercises of abilities can be multiply realizable. For instance, there may be situations in which one focusses one's binoculars with the help of an automatic focussing machine, which takes readings of one's eyes and adjusts the binoculars' lenses accordingly. (Indeed, the exercises in question might not even *most typically* be explained in this way. Especially if I am going to a familiar restaurant, ordering the right-size pizza may be best explained by appeal to some inductive judgement, such as 'I always order the medium pizza and that's worked out well every time.') But as long as they admit of sometimes being best explained by appeal to good introspective judgements, where these explanations are instances of the ordinary explanatory pattern I am drawing attention to, they count as introspection-reliant.

With the notion of introspection-reliant abilities in hand, we can turn to my proposal for when it is rational to hold given introspective judgements to be good. This is where we shift focus from the exercise of an ability back to its possession. The proposal follows the shape of the one developed earlier for visual experiences:

We can rationally regard introspective judgements as good when doing so helps explain the possession of environment-indexed introspection-reliant abilities.

Attributing good introspective judgements to a subject allows us to understand her having a range of abilities which are indexed to her current situation. In particular, it allows us to understand how she has it in her power to do certain things in the circumstances in which she finds herself (provided, of course, other factors are in place). This will be true independently of whether she in fact does any of those things in that situation. In this way, these attributions form part of common sense explanatory patterns: we habitually attribute judgements or other kinds of mental states to people so that we have a better handle on what kinds of agents we are. Part of this involves understanding what we are able to do in the situations we are in.

This is easy to see in Focusing Binoculars. The analogy with the visual case is tight here because focussing binoculars is a practical ability akin to the visual-based abilities mentioned earlier. Suppose you are sitting in tall reeds looking through your binoculars. You are poised to catch a glimpse of a rare Spotted Sandpiper, believed to be foraging somewhere in the distance. While peering through the lenses, you introspectively judge whether things look sharp to you. According to my proposal, we should treat your introspective judgement as good if doing so helps explain your possession of some introspection-reliant abilities—for instance, your ability *in situ* to focus the binoculars on the bird in the distance, your ability *in situ* to find out whether there is a Spotted Sandpiper foraging around over there by looking, and so on. In this case, appeal to good introspective judgements about the fuzziness of how things look, for instance, does allow us to understand how you are able, then and there while sitting in the reeds, to focus your binoculars on a bird in the distance, to locate the Spotted Sandpiper nearby, and so on—even if you do not do any of these things at that moment. Your good introspective judgement provides the appropriate sort of connection between you and aspects of the particular circumstances you are in, a connection which would be exploited were you to exercise an introspection-reliant ability, such as focussing your binoculars on the bird in front of you in your current situation. Good introspective judgement thus helps explain how you have it in your power to focus your binoculars, locate a nearby bird, etc in the situation you are in.

Consider Ordering Pizza: you are sitting in a restaurant about to order. While deliberating about what to do, you make introspective judgements about how hungry you feel and about what you would like to eat. We should treat your introspective judgements as good if doing so helps explain your possession of some abilities—for instance, your ability to order the kind of food (pizza or cannelloni) that will make for a satisfying eating event for you, or your ability to order an amount of food to match your hunger on this occasion. In such a case, appeal to good introspective judgements about what you like to eat, or how hungry you feel, allows us to understand how you are able, *in situ*, to order food that is apt in both its kind and amount—even if you do not in fact order anything at all in the end.

Lastly, consider Slowing Down. Suppose you have just printed out the papers for the Dean and looked up some applications to discuss with the admissions chair. You are poised to leave the office, key in hand, feeling very hassled and pressured. As you are considering all the things you have to do, how important it is to get them done, how little time you have available, and what the options are, you introspectively judge that you are thinking carefully—i.e. slowly and systematically about which options are available. We should treat your introspective judgement as good if doing so helps explain your possession of some abilities—for instance, your ability in this situation to slow down your reasoning involved in figuring out how to get everything done to ensure careful planning. Again, appeal to good introspective judgements about whether you are thinking things through properly and calmly allows us to understand how you are able to shift yourself from being somewhat

discombobulated into a position to carry out effective planning in high-stakes, high-pressure situations—even if you do not end up forming any plan at all.

All three cases involve abilities the exercise of which involves success at doing something with respect to a specific aspect of one's current situation. That success, or potential for success, is explained in each case by safe cognitive access to relevant aspects of one's environment, where this includes one's own mental condition. A subject's environment-indexed ability to focus her binoculars on the bird, order the right kind or amount of food, to slow down her thinking for effective planning under pressure, is explained by appeal to knowledge of looks or of hunger feelings, or of whether one is thinking slowly and systematically enough. These sorts of facts are typically, or sometimes, known by introspection.

Before finishing this section, I want to address some concerns about the examples I have given, concentrating on Ordering Pizza. The discussion applies, *mutatis mutandis*, to Focussing Binoculars and Slowing Down as well.

As I mentioned above, it's plausible that the decision about what to order and how much to order can also be made differently than it is made in Ordering Pizza. For example, it may be entirely based on induction from previous restaurant visits. Perhaps this is even how one decides in most cases. Moreover, it is also plausible to think that good ordering takes account of past successes and failures, so even when the decision isn't made entirely on inductive grounds it will involve doing so to some degree. But Ordering Pizza merely requires that there are *sometimes* cases where the following is a good explanation: we are in a position to order the right-size pizza at least *in part* because we know how hungry we feel right then, because we can gauge our hunger. The thought is that in such circumstances this knowledge typically is got via introspection—in that situation we introspectively judge how hungry we feel and we come to know this thereby.

One might argue that it is never a good explanation to explain possession of the ability to order the right-size pizza by appeal to introspective knowledge of how hungry we feel. One could do so for several different reasons. Firstly, one might object to the claim that knowledge of how hungry we feel, whether introspective or not, explains how we are in a position to make successful ordering decisions about how much and what to order. According to this objection, our ability to order the right-size pizza just does not get explained ever by appeal to knowledge of how hungry we feel. If this is right then introspective knowledge of how hungry we feel, even if we have it, plays no role here. I think that this is not a very plausible line of objection, given how minimal the demands of the case are. The objection involves maintaining that there is no significant correlation between our feelings of hunger and how hungry we in fact are, so that the former cannot be exploited at all as a gauge of the latter. That seems an unduly pessimistic view about how our hunger sensation corresponds to actual hunger states. It is worth remembering here that I do not claim that the ability to order appropriately is very fine-grained.

Secondly, although agreeing that knowledge of how hungry we feel sometimes can help explain having the ability to make a successful ordering decision, one might object to the claim that this knowledge comes via introspection. The suggestion is that insofar as we make

good judgements of how hungry we feel, they are not introspection-based.⁸ This objection is rooted in the more specific denial that relevant good judgements about how hungry we feel are introspective judgements. This is a genuine matter of debate and it cannot be settled by appeal to introspection, of course. There are different views about what counts as introspective judgement, and what sort of mental aspects (or even non-mental aspects) we can know by introspection. I do rely on some pre-theoretic way of being able to identify when we are introspecting and when we don't—just as, I think, we are all fairly happy to rely on such an understanding of what counts as a visual perception, say. There are borderline cases and opinions diverge about those. Some will think that first-personal judgements about how hungry we feel are introspective; some will think that they are akin to proprioceptive judgement, where the latter in turn are kinds of perceptions; some will think something else entirely. As I alluded to above, it is not so important for me that this particular example is convincing to everyone—I am happy to shift examples. It would only be worrisome for me if there was *no* instance where it was plausible that attribution of introspective knowledge helps explain possession of abilities—i.e. if someone made a case that introspective judgements played no role in our environment-indexed agency at all. That is, it would be worrisome if for any situation in which an introspective judgment can be attributed, it is idle with respect to our abilities. Thus, as long as there is a concession that there are some introspective judgements which put us in a position to something, to exercise some ability, this does not matter to the general proposal I defend here. I said at the beginning, that proposal is independent of a particular conception of introspection. To see it at work one must provide examples, though, and such examples do involve more specific views about which kinds of judgements count as introspective.

The proposal thus yields a reason-driven basis for moderation in our scepticism about introspection. Introspection can be rationally held to yield knowledge in cases where introspective judgements play a role in explaining possession of introspection-reliant abilities. On the assumption that we have some introspection-reliant abilities, then, we can accept that some of our introspective judgements are good.

4 Anchoring Common Sense Belief in Introspective Knowledge

The assumption that we have some introspection-reliant abilities is not extravagant, since many of the abilities in question are humdrum abilities that let us get on with run of the mill tasks in everyday life. We have a robust sense of which ordinary kinds of introspection-reliant abilities we possess. That robust sense of which introspection-reliant abilities we possess helps to rationally anchor common sense conviction about introspective knowledge. Let me explain this now.

The strength of our common sense conviction that we have introspective knowledge is manifested in the fact that ordinary introspective judgements are typically made with high

⁸ One version of this objection says that these judgements are not introspection-based because there is no psychological kind sensibly thought of as 'introspection'. In this paper, I have nothing to say to someone who does not accept that there are any introspective judgements at all. I start with the assumption that we can and do sometimes introspect and that the upshot of this is introspective judgement. My overall question is whether we have good reason to hold that we can know anything by introspection (and whether this knowledge can be exploited by scientific investigation), not whether there is anything that qualifies as introspection in the first place. This is a good issue, though, so I do not mean to dismiss debate about it. For example, one might worry that we cannot properly untangle introspective from non-introspective aspects of self-attributions (see Schwitzgebel 2012 for discussion of the difficulty to unpick introspective from non-introspective states).

confidence. What is more, this high confidence tends to be resilient. Suppose you introspectively judge that you have a headache. Your friend, upon looking searchingly at your face, tells you that you don't have your 'headache face' and so you probably don't have a headache. In response, you will not be likely to change your mind but you will continue to be highly confident that you have a headache. And so it would be with a host of other bits of superficial information which cut against your introspective judgement. In an ordinary case like this, you would likely stick to your high confidence in light of these pieces of information. These two aspects—high confidence and resilience—reflect the fact that common sense takes introspection to be a secure source of information. On the view I propose, this commitment to introspection as a secure source of information is itself rooted in corrective feedback concerning which abilities we possess and how we possess them. The key point is that the typical high confidence and resilience of introspective judgement is not only intuitive, but it is empirically warranted by such corrective feedback. I will spell this out further in just a moment, but first I want to emphasize that the view I am putting forward here flips the standard sceptical position on its head. According to the standard sceptical line, ordinary high confidence and resilience of introspective judgement is a sign of epistemic pathology. The high confidence of introspection has been claimed to be inappropriate, unwarranted and symptomatic of a lack of corrective feedback. It has been suggested, for instance, that in absence of any proper checks on introspective judgement we become overconfident, like a boastful person who is never told off (Schwitzgebel 2011: 130). I disagree with this diagnosis: our everyday practice of explaining the possession of environment-indexed abilities by appeal to good introspective judgement, itself makes for a great deal of corrective feedback. The high resilient confidence associated with ordinary introspective judgement is anchored in it. As we shall see, these aspects of ordinary introspective judgement are signs of their empirical well-groundedness, rather than their epistemic pathology.

So, what is this corrective feedback relevant to introspective judgements? Consider first that we can and do assess our abilities in everyday contexts, including our introspection-reliant ones. Day-to-day assessment of how well we can exercise our abilities is honed by awareness of success and failure of such exercise in a variety of situations. In turn, this responsiveness allows us also to fine-tune some of our abilities—we get better at doing things. Our robust sense of our ordinary abilities is therefore empirically grounded in this everyday way. But such empirically-informed conviction in our abilities encourages corresponding strength and resilience of introspection-based confidence by means of which we typically exercise those abilities. Consider the following analogy. Suppose you start hearing a voice telling you every morning whether the bus you take to work is going to be late. The prediction made by the voice turns out to be right every time (or close to every time) about the busses. After a while, you are going to start trusting that voice and guiding your mornings accordingly. The more feedback you have about the voice getting things right—e.g. the more it enables you to have stress-free and efficiently-timed mornings—the higher your confidence in the voice's predictions and the more resilient that confidence will become. If following the voice's forecast has resulted in relaxed breakfasts and commutes to work for several months, then, despite the fact that this morning's radio news announcement says that busses are running on time, when the voice predicts that your bus is going to be late you are still going to be confident in that prediction and pour yourself another cup of tea. In a similar vein, one's confidence in possessing a certain ability, where that confidence is grounded in plenty of feedback from situations in which one has exercised or attempted to exercise the ability, also promotes resilient high confidence in states in virtue of which we have that ability. Insofar as

such high confidence manifests common sense conviction that we have introspective knowledge, that conviction seems a perfectly reasonable take on the situation.

In section 1, I briefly mentioned that one important strand of scepticism about the use of introspective data in theorizing focuses on the privacy of such data. The worry is that this privacy precludes corrective feedback from other measuring devices with which to check the accuracy of introspection. But, inter-measurement agreement, which requires publicly available introspective data, is not the only way to calibrate introspection. As I have argued just now, we have a great deal of empirical corrective feedback on the accuracy of introspection via its explanatory role vis-a-vis our environment-indexed abilities. That feedback, we have seen, grounds common sense commitment to introspective knowledge. Like any other *a posteriori* empirically-grounded commitment, moreover, it is subject to systematic empirical investigation and revision. So one major advantage of the view defended here is that it is not merely compatible with such empirical investigation, but it positively encourages engagement with relevant empirical work. Let me explain this.

It is true that the starting point of the proposal I put forward is our common sense explanatory practice and the relevant common sense view of our abilities reflected in it. But, as I point out above, common sense in these matters is not from the armchair in the first place—the explanatory practice incorporating views about our ordinary abilities is responsive to empirical feedback in everyday life. This is because our ordinary conceptions of our abilities and of the extent to which we possess them are shaped by constant feedback from exercising our abilities or observing people exercising abilities over time. This empirical input can take a more rigorous and deliberate turn in the hands of a psychologist. In particular, there are empirical investigations of *both* the relevant abilities (possession of which is explained by appeal to introspective knowledge, according to my proposal), and our ordinary grip on the extent to which we possess these abilities.

For instance, empirical studies show that everyday assessment sometimes significantly and systematically over- or under-estimates some abilities. There are detailed empirical studies of different abilities and our common conceptions of them. These studies show that sometimes the two are significantly out of line with one another. A well-known example of this involve our ability to provide accurate eyewitness testimony (e.g. Loftus 1979; Lindsay 1994; Devenport *et al.* 1997; Wells & Loftus 2003). Often, the misconception is due to implicit associations or unconscious bias as for instance in the case of our ability to evaluate and interview applicants in a gender-neutral way (e.g. Schein 1973; Harris 1989; Steinpreis *et al.* 1999; Goldin & Rouse 2000; Eagly & Karau 2002). These studies make clear that our everyday grip on our abilities is fallible. They also make clear, more importantly for our present purposes, that our everyday grip on our abilities is open to revision in light of empirical investigation.

This is good news for my proposal: in addition to the basic empirical grounding of our everyday grip on abilities which comes from responsiveness to success and failure of their exercise in ordinary settings, we can also improve that grip with empirical data from scientific study. Thus we have within reach a sophisticated and scientifically-respectable constraint on verdicts concerning goodness or otherwise of introspective judgement.

What is more, my proposal can explain how ordinary intuition concerning introspective knowledge itself, grounded as it is in common sense explanatory practice, could eventually be replaced by a more sophisticated, scientifically grounded intuition about introspective

knowledge. In general, we have a picture of our folk understanding of the world as becoming gradually and to varying extents informed by increasing scientific understanding of the world. Common sense often absorbs the advancements of science. For instance, it used to be part of common sense that the earth and humans are the centre of the universe, that there are no sub-conscious mental states, and that solid objects (such as rocks and trees) occupy every bit of space in the overall region they occupy. None of these assumptions are part of common sense now since it has been influenced by developments in astronomy and microphysics, and by Freudian psychology. In a similar manner, increasingly widespread scientific knowledge about our abilities and about our systematic over- and underestimation of them can lead both to revision and correction of our common sense understanding of these matters. In turn, a revised conception of our abilities would lead to a revised intuition about the kind of cognitive access we have in terms of which we explain the having of such abilities in a given situation. My proposal therefore suggests a particular mechanism for the interaction between common sense intuition and scientific feedback, with the mechanism being a of a sort that we already accept in general.

5 How to Solve the Calibration Problem

I have put forward a proposal about the rational basis of common sense commitment to introspective knowledge. This proposal delivers a strategy for solving the calibration problem for introspection. Indeed, we can make progress with both dimensions of the calibration task mentioned at the beginning of this paper in the passage cited from Goldman. We can get a fix on the kind of conditions under which introspection tends to produce good introspective judgements ('operational conditions'), and on the kinds of contents for which introspective judgements tend to be good ('content conditions'). Our empirically-informed understanding of our ordinary abilities allows us—as theoreticians—to pin down both kinds of conditions.

Consider operational conditions first. The fact that we can become aware of the relative success or failure of an attempt to exercise an introspection-reliant ability in different situations offers a way to specify optimal conditions for its exercise. When considering several different introspection-reliant abilities, certain operational conditions will stick out as characteristic of abilities that involve introspection. Given a large enough sample, conditions in this group will stand out as something like optimal conditions for introspection to produce good introspective judgments. Being able to order the right-size pizza in a given situation requires (among other things) that I am alert, not distracted, not under the influence of drugs, etc. So does slowing down one's thinking under pressure. And so does focussing binoculars when hiding in the reeds. All these environment-indexed abilities have in common that their possession typically depends on good information got by introspection and successful exercise of each of them typically depends on certain general conditions obtaining. This makes it plausible that evidence concerning optimal conditions for such abilities yields evidence concerning optimal—or at least favourable—conditions for introspection. These will be the operational conditions for introspection.

Consider content conditions next. Two subjects can have the same broad type of ability even though one of them is better at exercising it. Nina Simone was able to sing 'I put a spell on you' and so am I, but Nina Simone could do it much better. Sometimes the difference in how well two people are able to exercise a given ability is due to operational conditions. Perhaps Nina Simone sang better than I do because I always try to sing in sub-optimal conditions. But other times the difference is due to a more systemic limitation, where this invokes a difference in possession of finer-grained abilities. For example, try as you may in optimal

conditions, you will never be able to keep up with a cheetah running flat out, even though you and it are both able to run fast. Humans and cheetahs have different kinds of bodies, and some of their differing physiological and mechanical properties account for their differing abilities for speed when running flat out. Consider another example: an eagle and I both have the ability to visually track objects at a distance. But an eagle can track the movements of a dormouse from a distance at which I cannot do so, even if my eyesight and the lighting are perfect. The reason is that eagles have a more extended range of visual acuity than humans. Their visual experiences present things at a higher resolution. The eagle and I are both able to track objects at a distance, but the eagle can have visual experiences with contents that I cannot have, and so the eagle can do better than I can at tracking things. The difference in our respective visual contents explains why, in certain circumstances, the eagle has a finer-grained ability to track objects at a distance than I have. This example illustrates that one significant factor in the explanation of possession of certain abilities concerns content-bearing states. A key explanatorily efficacious feature is the possession of specific contents.

The content of introspective judgements, too—not merely their goodness—contributes crucially to their explanatory efficacy with respect to possession of many introspection-reliant abilities. While I am waiting in the reeds, for example, my introspective judgement that I have a headache, even if good, does nothing to explain my possession of the ability to focus my binoculars on the bird in the distance. In order for an introspective judgement to play that role it must have relevant content: it must be useful, in light of my background beliefs, with respect to focussing my binoculars on the bird. Most directly, it will be useful if it contains fairly detailed information about the sharpness or fuzziness of how things look. Similarly, when facing the waiter in a restaurant, not just any good introspective judgment will explain my ability to order the right amount of food in such a situation. Introspective judgement must provide information about how hungry I feel if it is to do so, and this is a constraint on the content of that introspective judgement.

In addition to demanding that they are good judgements, the explanatory role of introspective judgements places constraints on their content. Via these explanatory demands we can pin down types of content for which introspection tends to produce good judgements. But how strong a constraint is this? We are interested in whether it is substantial enough to be employed in disputes about the legitimacy of introspective evidence in various areas of theorizing about the mind. I think the constraint can do significant work in this respect. We can use it to home in on content not only in terms of rough subject-matter, but also go some way to fixing levels of detail or resolution. Let me close with a brief sketch of how to apply the constraint.

We tend to be very successful at focussing binoculars—binoculars normally do let us find out about and track distal objects. In typical cases we can explain this success by appeal to good introspective judgements about how fuzzy or sharp things look to a fairly high degree of resolution. This explanation does not require, though, attributing good introspective judgements the contents of which might settle controversial question about visual experience: questions about whether visual fuzziness is a matter of representational or non-representational aspects of the experience, questions about the exact size of the area of the visual field where things are presented stably and sharply (as opposed to any more fuzzy regions at the periphery), and so on.⁹ Contents that can settle such questions make no

⁹ For discussion of whether some instances of visual blurriness are a matter of representational or non-representational properties of visual experience see e.g. Boghossian and Velleman 1989, Tye 2000, Dretske

difference to the explanatory efficacy of introspective judgement with respect to abilities of the Focussing Binoculars type.

We are also pretty decent at ordering the right kind and amount of food in restaurants. However, it is not uncommon to realize that one has over- or under-estimated one's hunger, resulting in a bite left on the plate or a residual desire for just a little more. This suggests that on average we have the ability to successfully order in-the-ballpark-right amounts of food, but not the ability to order correct amounts down to the exact morsel. In typical cases we can explain possession of the ability we do have by appeal to good introspective judgements about whether one is very hungry, moderately hungry, or not very hungry (or something like that). Such an explanation does not require attributing good introspective judgements with contents about hyper-precise levels of hunger. When we consider introspection-reliant abilities such as ordering the right amount of food, then, possession of which can be explained by appeal to good introspective judgements about how hungry one feels, introspection tends to be accurate for contents containing fairly coarse characterisations of one's hunger.

Some of us—the proverbial ‘cooler heads’—are also fairly good at calming ourselves down to ensure difficult planning when under pressure. We can often get ourselves into an improved position to figure out how best to get done a number of high-stakes tasks in a very limited amount of time even if it is not obvious how at first glance. Having such an ability can be, typically or sometimes, explained by appeal to good introspective judgements about whether one is thinking carefully—slowly, deliberately, systematically—about the issues at hand. These are introspective judgements about aspects of our thinking which render it conducive to efficient and successful planning. We can tell introspectively when we are thinking in a manner favourable to complicated planning but also when we are not. As I suggested above, in the latter case, we can exercise some limited control over our thinking to ensure that we are thinking things through properly. When feeling on the verge of panicking, we can slow down our stream of thought and exert ourselves to proceed more methodically in our deliberations. In light of such abilities, possession of which can be explained by appeal to good introspective judgements about (e.g.) how fast and systematically one is thinking, introspection tends to be accurate for contents containing certain markers of effective deliberation and systematicity of thinking. What these contents are more precisely, which markers we are tracking introspectively, is subject to further empirical discovery. It certainly does not fall out of simple introspective reflection.

Introspection tends to be epistemically successful for types of content which are, as I think of them, middle-of-the-road. They are vastly more detailed and informative than the frugal judgements which serious sceptics grudgingly admit figure into introspective knowledge—e.g. ‘red phenomenology now’, ‘vivid pain now’ (Schwitzgebel 2011: 129). But they provide much less information than the theoretical types of judgements which are used in various debates about the nature of consciousness. The latter kinds of judgements are, of course, central targets of sceptical arguments in the literature—such as ‘visual experience is clear and

2003. For discussion of the relative size and difference of the central foveal and peripheral regions of the visual field, see Dennett 1991, Schwitzgebel 2011.

stable one hundred degrees into the periphery’, ‘emotional experience is entirely bodily defined’, or ‘conscious thought has a *sui generis* phenomenal character’.¹⁰

But this seems the right result. Ordinarily we take for granted that introspective judgements such as ‘I have a headache’, ‘I like burritos’, or ‘I’m not seeing things clearly’ can amount to knowledge, and we are pretty confident about that. These are middle-of-the-road kinds of judgements. It is overwhelmingly plausible, given their explanatory role vis-a-vis our environment-indexed abilities, that introspection tends to produce knowledgeable judgements of their kind.

Can such middle-of-the-road introspective knowledge be useful to philosophy or to science? There is certainly room for optimism here. The strategy for calibrating introspection I have sketched shows that there are non-trivial operational and content conditions for introspection. We have also seen how these conditions are subject to refinement and revision in light of empirical results. We should thus expect calibrated introspection to deliver non-trivial and scientifically respectable data about conscious experience. These introspective data should be highly relevant to scientific and philosophical investigations of consciousness.¹¹

References

- Bayne, T. & Spener, M. 2010. ‘Introspective Humility.’ In *Philosophical Issues* 20 (1): 1–22.
- Boghossian, P. & Velleman, D. 1989. ‘Colour as a Secondary Quality.’ *Mind* 98: 81–103.
- Byrne, A. 2011. ‘Knowing That I am Thinking.’ In *Self-Knowledge*, A. Hatzimoysis (ed), Oxford University Press.
- Cummins, R. 1975. ‘Functional Analysis.’ *Journal of Philosophy* 72: 741–765.
- Cummins, R. 1983. *The Nature of Psychological Explanation*. MIT Press.
- Cummins, R. 2000. “‘How Does It Work?’” vs. “‘What Are the Laws?’” Two Conceptions of Psychological Explanation.’ In F. Keil & R. Wilson (eds), *Explanation and Cognition*, MIT Press.
- Dennett, D. 1991. *Consciousness Explained*. Penguin Books.
- Dretske, F. 2003. ‘Experience as Representation.’ *Philosophical Issues: Philosophy of Mind* 13: 67–82.

¹⁰ My point here is not that questions about, say, the size of the clear and stable foveal centre in the visual field, or other theoretical aspects of conscious experience are bad questions to ask. They are bad questions to ask *introspection* to settle.

¹¹ I thank John Bengson, David Chalmers, Cian Dorr, Anthony Eagle, Miguel Fernandez, Anandi Hattiangadi, Ian Phillips, Hannah Pickard, Daniel Rothschild, Nick Shea, Susanna Siegel and Scott Sturgeon for helpful comments on various drafts of this paper.

Devenport, J.L., Penrod, S. D. & Cutler, B. L. 1997. 'Eyewitness Identification Evidence: Evaluating Commonsense Evaluations.' *Psychology, Public Policy and Law* 3: 338–361.

Eagly, A.H. & Karau, S.J. 2002. 'Role Congruity Theory of Prejudice Toward Female Leaders.' *Psychological Review* 109, no. 3, pp. 573–598.

Gertler, B. 2010. *Self-Knowledge*. Routledge.

Gertler, B. 2011. 'Self-Knowledge.' *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/self-knowledge/>.

Goldin, C. & Rouse, C. 2000. 'Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians.' *American Economic Review* 90 (4): 715–741.

Goldman, A. 2004. 'Epistemology and the Evidential Status of Introspective Reports.' In *Trusting the Subject? Vol 2*, A. Jack & A. Roepstorff (eds) Imprint Academic.

Harris, M. M. 1989. 'Reconsidering the Employment Interview: A Review of Recent Literature and Suggestions for Future Research.' *Personnel Psychology* 42: 691–726.

Lindsay, D. S. 1994. 'Memory Source Monitoring and Eyewitness Testimony.' In *Adult Eyewitness Testimony*, D. Ross, J. Read, and M. Toglia (eds). Cambridge University Press.

Loftus, E. F. 1979. *Eyewitness Testimony*. Harvard University Press.

Schein, V.E. 1973. 'The Relationship Between Sex Role Stereotypes and Requisite Management Characteristics.' *Journal of Applied Psychology* 57: 95–100.

Schwitzgebel, E. 2008. 'The Unreliability of Naive Introspection.' *Philosophical Review* 117 (2): 245–273.

Schwitzgebel, E.. 2011. *The Perplexities of Consciousness*. MIT Press.

Schwitzgebel, E.. 2012. 'Introspection, What?' In *Introspection and Consciousness*, D. Smithies & D. Stoljar (eds.), Oxford University Press.

Shoemaker, S. 1994. 'Self-knowledge and "Inner Sense"'. *Philosophy and Phenomenological Research* 54: 249–314.

Spener, M. 2011a. 'Disagreement about Cognitive Phenomenology.' In *Cognitive Phenomenology*, T. Bayne & M. Montague (eds.), Oxford University Press.

Spener, M. 2011b. 'Mind-Independence and Visual Phenomenology.' In *Introspection and Consciousness*, D. Smithies & D. Stoljar (eds.), Oxford University Press.

Spener, M. Manuscript. 'Phenomenal Adequacy and Introspective Evidence'.

Steinpreis, R. E., Anders, K. A., & Ritzke, D. 1999. 'The Impact of Gender on the Review of the Curricula Vitae of Job Applicants and Tenure Candidates: A National Empirical Study.' *Sex Roles* 41: 509–528.

Tye, M. 2000. *Consciousness, Color, and Content*. MIT Press.

Wells, G. L. & Loftus, E. F. 2003. Eyewitness Memory for People and Events. In Handbook of Psychology: Forensic Psychology, vol 11. A. Goldstein (ed). Wiley.