

# Philosophical and Psychological Accounts of Expertise and Experts

*Matt Stichter*<sup>†</sup>  
mstichter@wsu.edu

## ABSTRACT

There are many philosophical problems surrounding experts, given the power and status accorded to them in society. We think that what makes someone an expert is having expertise in some skill domain. But what does expertise consist in, and how closely related is expertise to the notion of an expert? In this paper I inquire into the nature of expertise, by drawing on recent psychological research on skill acquisition and expert performance. In addition, I connect this research on expertise to the larger context of psychological research on human cognition, as it will illuminate some of the differing elements of expertise. This allows me to then critique philosophical accounts of expertise, by showing how they make unwarranted assumptions about skills and expertise. Finally, I note the ways in which being credited as an expert can diverge from the possession of expertise itself. This can help us resist some of the power dynamics involved with those deemed to be experts.

Keywords: dual-process, expertise, philosophy, psychology, skill.

## Introduction

There are many philosophical problems surrounding experts, given the power and status accorded to them in society. We think that what makes someone an expert is having expertise in some skill domain.<sup>1</sup> But what does expertise consist in, and how closely related is expertise to the notion of an expert? Although most of us have acquired several practical skills, few of us have achieved the level of expertise with regard to those skills. So we can be easily misled as to the nature of expertise, since it differs significantly from earlier

<sup>†</sup> School of Politics, Philosophy, & Public Affairs, Washington State University, US.

<sup>1</sup> Expertise refers to the highest level of skill acquisition, for the possession of a skill is a matter of degree.

stages of skill acquisition. Furthermore, this potential for misleading characterizations of skills and expertise leads to philosophers implicitly working with different conceptions of skills. This can interfere with their attempts to solve related problems about experts.

In this paper I inquire into the nature of expertise, by drawing on recent psychological research on skill acquisition and expert performance.<sup>2</sup> In addition, I connect this research on expertise to the larger context of psychological research on human cognition, as it will illuminate some of the differing elements of expertise. This allows me to then critique philosophical accounts of expertise, by showing how they make unwarranted assumptions about skills and expertise. Finally, I note the ways in which being credited as an expert can diverge from the possession of expertise itself. This can help us resist some of the power dynamics involved with those deemed to be experts.

The first section of this paper provides an overview of the central features of expert performance from the perspective of the current psychological research on expertise. The main two features are automaticity and a recognition-primed form of decision making. Following this is a discussion of the implications of these features of expertise for the articulation and codification of expert knowledge. This section also briefly covers the distinction between System 1 (intuitive) and System 2 (deliberate) thinking in cognitive science, and how aspects of expertise draw on both systems. The second section focuses on how expertise is acquired. The main two features are deliberate practice and self-regulation. Following this is a discussion of the implications of the features for the role of motivation in acquiring expertise. The third section of this paper provides an overview of the Dreyfus model of expertise, and the fourth covers Julia Annas's account of expertise. Both of these philosophical accounts are critiqued from the standpoint of the recent psychological research on expertise. The final section brings out important distinctions between having expertise and being credited as an expert.

## 1. Expert Performance

Expertise can be thought of in two ways: with respect to a specific skill; or with respect to a domain, where expertise is a collection of related skills.<sup>3</sup> In either

<sup>2</sup> Hereafter, when I mention “research” I’ll be referring to the psychological research on expertise.

<sup>3</sup> A skill can roughly be defined as a learned ability to achieve a desired outcome, though often it goes undefined even in the psychological literature. It’s important to note that a skill involves some flexibility in

case, a description of expertise often proceeds by comparing the performances of experts with novices. A defining feature of expert performance is the ability of experts to act in a way that seems (and usually is) almost effortless. Experts do not need to devote much conscious attention to what they are doing, and this lack of conscious attention does not lead to any reduction in their performance. This phenomenon is referred to as automaticity in the psychological literature.<sup>4</sup> While automaticity is a defining feature of expert performance, it starts to appear at earlier stages of skill development. With practice, tasks can be accomplished more effectively and more efficiently. This allows a person to devote less attention to the tasks at hand without any reduction in performance, and to shift that attention to other matters. Being able to improve one's performance requires having the initial tasks becoming effortless, so one can devote attention and energy to more difficult tasks. This highlights the importance of Csikszentmihalyi's (1993) work on flow, where one is fully immersed in the task at hand. Being in this state means that you do not need to exert self-control to keep yourself from being distracted. Not only does flow free up one's attention, but it also makes one more unlikely to be disrupted by external distractions.

Another way in which automaticity enables effortless expert performance is by allowing the expert to operate well on the basis of intuitive (rather than deliberative) judgments, as intuitions are experienced as immediate and not as the result of any conscious deliberation. This intuitiveness is central to expert performance because it allows the expert to react quickly to situations. One important thing to keep in mind about the talk of intuition in expertise is that the ability of the expert to reliably act well on an intuitive level is due to having an immense amount of experience and practice.<sup>5</sup>

Expertise, however, is not the only source of intuitive judgment. Intuitions can also arise from the use of mental heuristics, which are basically short-cuts in reasoning, where you simplify a complex problem in order to come to a decision more easily. Since there are multiple sources of intuitive judgments, and they vary with respect to reliability, it will be important to cover a

how one goes about achieving that outcome (to cope with changes in one's environment), as well as a broad view of the outcome (as in learning how to speak a language, rather than a single phrase).

<sup>4</sup> Feltovich, Prieluda and Ericsson, 2006, p. 53

<sup>5</sup> The psychological research "locates automaticity on the backend of development. It is the outcome of repeated experience, of instruction, intentional coaching and socialisation." (Lapsley and Hill, 2008, pp. 324-325)

distinction between two general types of cognitive processing: System 1 and System 2.<sup>6</sup> The distinction between System 1 and System 2 thinking is now fairly commonplace thanks to the work of the psychologist Daniel Kahneman. The distinction between these two Systems is important for understanding expertise, especially the intuitive side of it. System 1 refers to the intuitive side of our mental life, which is automatic and spontaneous. System 2 refers to the kind of thinking we identify with agency – making deliberate choices between options, engaging in various forms of analysis, and exerting self-control. Generally, our behavior is guided by System 1, unless we choose to stop and think about what we’re doing, in which case System 2 takes charge. While it might sound from this that we ought to let System 2 take the reins most of the time, that turns out not to be the case. Kahneman’s work shows that most of the time System 1 guides us efficiently and effectively, when you consider that by default most of our actions are guided by it. In addition, System 2 requires deliberate effort and attention, which is mentally taxing, and so it limits how often we can engage this system.

Kanheman’s work focuses on the heuristics that are used in System 1 thinking to generate intuitive solutions to problems we encounter, especially the unreliability of heuristics. While heuristics provide us with good solutions in many circumstances, they are also the source of systematic biases or errors. The “availability heuristic”, for example, is used when people judge the probability of an event occurring based on how easy it is to recall examples of those events. While an event being frequently mentioned might be due to it occurring often, it’s also the case that more sensational events get mentioned more frequently (like shark attacks on swimmers which occurs rarely). What’s most important for the present purposes, however, is to note that heuristics are not the sole source of intuitive judgments. As Kahneman himself notes:

the accurate intuitions of experts are better explained by the effects of prolonged practice than by heuristics. We can now draw a richer and more balanced picture, in which skill and heuristics are alternative sources of intuitive judgments and choices. (Kahneman, 2011, p. 11).

Skills are context sensitive, and the accuracy of the intuitive judgments that arise in expertise is due to the great familiarity the expert has in operating in these kinds of situations. For example, the chess expert can have a reliable

<sup>6</sup> This is also known as dual-processing theory. For a critique of this theory, see Kren, G. and Schul, Y. (2009). For a defense from this critique, see Evans, J. and Stanovich, K. (2013).

intuition about what move to make in a situation because of her familiarity with being faced with this kind of board position before. The grounding of intuitions in this way is the reason why the intuitive judgments of experts are seen as highly reliable, in contrast to the inconsistent reliability of intuitive judgments produced by heuristics.

It is important to point out a limiting condition on the development of reliable intuitions arising out of experience and practice. Kahneman points out that we can't necessarily expect expertise to be achieved in all domains. He draws our attention to this in his overall description of what is required to develop accurate intuitive judgments:

The acquisition of skills requires a regular environment, and adequate opportunity to practice, and rapid and unequivocal feedback about the correctness of thoughts and actions. When these conditions are fulfilled, skill eventually develops, and the intuitive judgments and choices that quickly come to mind will mostly be accurate. (Kahneman, 2011, p. 416)

As noted in other accounts of skill acquisition, practice and feedback are essential.<sup>7</sup> But in order to get useful feedback when one practices, there needs to be some predictability in the environment itself, in the sense that “there are stable relationships between objectively identifiable cues and subsequent events or between cues and the outcomes of possible actions.”<sup>8</sup> Practice and feedback are what enable one to pick up on these cues at an intuitive level. However, if there isn't regularity between cues and subsequent events or outcomes, then recognition of those cues won't help you to figure out what to do next. For example, Kahneman argues that there doesn't seem to be enough regularity to the stock market environment to develop expertise in predicting stock prices.

Assuming there is enough predictability in one's environment to allow for the possibility of expertise; intuitive judgment can develop as you recognize cues from similar past experiences, and the outcome of actions that were taken in response. When you recognize that you have been in this situation before, and you have acted successfully in past situations like this one, then you do not

<sup>7</sup> “Whether professionals have a chance to develop intuitive expertise depends essentially on the quality and speed of feedback, as well as on sufficient opportunity to practice.” (Kahneman, 2011, p. 241) Kahneman, as an example, compares learning how to drive a car with learning how to pilot large ships in a harbor. The latter is more difficult to learn in part because of the longer delay between actions and noticeable consequences, which leads to slower feedback on one's attempt to pilot.

<sup>8</sup> Kahneman and Klein, 2009, p. 524

need to stop and deliberate about what to do next. This lack of deliberation is supported by the recognition-primed decision model, which was developed with extensive research on the decision making of fireground commanders.<sup>9</sup> This is clearly System 1 thinking at work. All of that practice and experience shapes experts such that:

experts see the world differently (Johnson and Mervis 1997; Myles-Worsley, Johnston, and Simons 1988). Because they have more and better organized knowledge in a domain, experts perceive things differently than do novices. They perceive different affordances. Perception of affordances is highly influenced by the amount of experience that one has with similar situations. (Lapsley, Narvaez, 2005, pp. 150-151).

Experience not only changes how experts view a situation, it also enables them to efficiently and effectively respond to the situation. A skilled chess player can know which moves to make because of her experiences in playing the game: being in a variety of situations, seeing the possible moves, and knowing which moves worked and which did not. What the player recognizes “includes the type of situation this is, what to expect from the situation (expectancies), suitable goals, typical courses of action (COAs), and relevant cues.”<sup>10</sup> This allows her to have an immediate intuitive response about what to do next in the situation.

Part of what follows from this is that unfamiliar or unusual situations will require the expert to deliberate to some extent about what to do, because the expert recognizes that the current situation doesn't easily map onto a previous situation. The expert can run a kind of mental simulation on the initial course of action [COA] that occurs to her. According to the recognition-primed model of decision making:

Mental simulation is the process of consciously envisioning a sequence of events, such as imagining how a COA [course of action] will play out. This allows a decision maker who knows enough to make accurate predictions to see what the consequences of a particular COA might be. . . . If the first COA

<sup>9</sup> Fireground commanders are those commanding firefighters on the scene of a fire. They have to arrive at decisions about how to coordinate the activities of the firefighters to contain the fire and keep everyone safe, based on the behavior of the fire and the skills of their firefighters (amongst other factors). “Data analysis found that approximately 80% of the commanders’ decisions were recognition-based. In fact, some interviewees said that they never made “decisions” at all.” (Klein, Ross and Shafter, 20006, p. 407).

<sup>10</sup> Klein, Ross and Shafer, 2006, p. 406

evaluated is found wanting, the expert generates a second and so on, evaluating each in turn but never comparing options against each other.<sup>11</sup>

Initially, one should not necessarily expect expert-level performance from an expert facing a unique situation, even if the experts will reliably perform better than non-experts in such situations. Expertise is limited to a certain background of experience.

Mental simulation is one of the places where we see the interplay between intuitive System 1 and deliberative System 2 in expert performance – System 1 provides the plan, and System 2 checks it.<sup>12</sup> It should be noted that while this kind of mental simulation is a form of System 2 thinking, it does not involve consciously comparing options or applying rules. Attempting to apply rules is what you do when you don't have any experience to draw upon. The courses of action that experts are simulating are drawn from their experience, and occur to them at an intuitive level. Furthermore, there's evidence "showing that when skilled decision makers abandoned their initial COA in favor of one they generated subsequently, the quality of that subsequent COA was significantly lower than their initial COA."<sup>13</sup> So while mental simulation can be useful for experts in some situations, it's not the case that engaging in System 2 thinking always improves upon the results.

While mental simulation engages System 2 thinking, it still operates on a course of action that was initially generated by System 1. So while experts might be able to articulate some of the process of mental simulation itself, they still cannot necessarily explain why they saw situations in a particular light, or why a particular course of action occurred to them.<sup>14</sup> The psychological research demonstrates that "experts often cannot articulate their knowledge because much of their knowledge is tacit and their overt intuitions can be flawed".<sup>15</sup> One reason for the difficulty in articulation is that intuitions arising out of expertise "are due to highly valid cues that the expert's System 1 has learned to use, even if System 2 has not learned to name them."<sup>16</sup> That is,

<sup>11</sup> Klein, Ross and Shafer, 2006, p. 406-407

<sup>12</sup> "The process involves both System 1 and System 2. In the first phase, a tentative plan comes to mind by an automatic function of associative memory – System 1. The next phase is a deliberate process in which the plan is mentally simulated to check if it will work – an operation of System 2." (Kahneman, 2011, p. 237)

<sup>13</sup> Klein, Ross and Shafer, 2006, p. 410

<sup>14</sup> Ericsson points out that "they cannot report why only one of several logically possible thoughts entered their attention, they must make inferences or confabulate answers to such questions." (Ericsson, 2006b, p. 230)

<sup>15</sup> Chi, 2006, p. 24

<sup>16</sup> Kahneman, 2011, p. 240

being asked to give an explanation of one's actions engages System 2 thinking. However, since the recognition of the situational cues and the resulting intuitive judgments are the work of System 1 thinking, an expert can't necessarily explain that part of her cognitive process.<sup>17</sup>

Even when experts are able to articulate an explanation, the explanations are often inconsistent with the observed behavior of the experts. These problems occur both when experts are asked about a specific task they just performed and when asked in general about their methods.<sup>18</sup> Of particular difficulty is getting an answer to the question of why the expert responded one way rather than another. It is important to note, however, that the research does not support the stronger conclusion that experts can never accurately articulate their reasons for action.<sup>19</sup> Rather, there are reasons why such articulation may be inherently difficult, and so articulation is not seen as a hallmark of expertise. In short, expertise is defined by performance, and such those with expertise have not been found to be able to reliably given accurate accounts of their decisions and judgments.

While there was an early hope in expertise research that the knowledge of experts could be extracted and rules could be developed that would greatly reduce the time it took to attain expertise, the problems with getting experts to articulate their knowledge reduced that hope.<sup>20</sup> Even if it was easier to get experts to articulate their knowledge, there's still a problem with trying to map

<sup>17</sup> Researchers involved with expert decision making maintain that "expert knowledge is largely tacit knowledge and can be difficult for the expert to share when asked. We cannot tell someone how to perform largely unconscious processes." (Klein, Ross and Shafer, 2006, p. 412)

<sup>18</sup> Ericsson notes that "When experts are asked to describe their general methods in professional activities, they sometimes have difficulties, and there is frequently poor correspondence between the behavior of computer programs (expert systems) implementing their described methods and their observed detailed behavior when presented with the same tasks and specific situations." (Ericsson, 2006b, p. 231)

<sup>19</sup> Despite these problems, there is a kind of reporting that experts can do about their thought process which does appear to be reliable. Instead of asking experts to explain their behavior after performing some task, experts are asked to 'think aloud' while engaged in performance of the task. While these verbalizations are far more accurate than after the fact explanations, they are not particularly detailed. The reason is that in 'think aloud' experiments "participants were not asked to describe or explain how they solve these problems and do not generate such descriptions or explanations. Instead, they are asked to stay focused on generating a solution to the problem and thus only give verbal expression to those thoughts that spontaneously emerge in attention during the generation of the solution." (Ericsson, 2006b, p. 228)

<sup>20</sup> Furthermore, the discovery of the complex of adaptations that mediate expert performance dispelled "the hope that it would be possible to extract the accumulated knowledge and rules of experts and then use this knowledge to more efficiently train future experts and, thus, reduce the decade or more of experience and training required for elite performance." (Feltovich, Prietula and Ericsson, 2006, p. 61) In addition, Polanyi is often credited as "the first critic who saw that nonconscious and intuitive mediation limits the possibility of eliciting and mapping the knowledge and rules that mediates experts' intuitive actions." (Ericsson, 2006a, p. 12)



out this knowledge, given the complexity of the mechanisms that mediate expert performance.

For example, Allen Newell (personal communication) described a project in which one of his graduate students in the 1970s tried to elicit all the relevant knowledge of a stamp collector. After some forty hours of interviews, Newell and his student gave up, as there was no sight of the end of the knowledge that the expert had acquired. As it may be difficult, perhaps impossible, to describe all the knowledge and skills of experts.<sup>21</sup>

It is important to keep realistic expectations of our ability to describe the knowledge of experts. Of course the research on expertise can extract some of the knowledge of experts, which helps to improve skill acquisition at all levels of performance, but there's no substitute for the role of experience in a variety of situations to achieve expertise.<sup>22</sup> This experience is what allows experts to reliably act well in an automatic and intuitive way.

## 2. Acquiring Expertise

While expert performance falls mainly within the domain of System 1 (intuitive) processing, there are a variety of ways in which System 2 (deliberate) comes into play in skill acquisition, which is what this section focuses on. Deliberate practice and self-regulating behavior are both essential parts of acquiring expertise, and these draw on System 2. Novices learning a skill will have to pay a lot of attention to what they are doing, and will need to expend a lot of deliberate effort in learning the basics of the skill. Since both self-control and cognitive effort draw on System 2 resources, and this is a limited pool of resources, the more you have to exert self-control to stay focused on the task at hand, the less cognitive effort you can expend on that task.<sup>23</sup> In order to make progress in learning a skill, the currently effortful tasks need to become effortless, in order to free up your attention to handle more complicated tasks.

<sup>21</sup> Ericsson, 2006b, pp. 235-236

<sup>22</sup> "All the paths to expert performance appear to require substantial extended effortful practice. Effortless mastery of expertise, magical bullets involving training machines, and dramatic shortcuts, are just myths. They cannot explain the acquisition of the mechanisms and adaptations that mediate skilled and expert performance." (Feltovich, Prietula and Ericsson, 2006, p. 61)

<sup>23</sup> Kahneman explains that you have "a limited budget of attention that you can allocate to activities, and if you try to go beyond your budget, you will fail. It is the mark of effortful activities that they interfere with each other, which is why it is difficult or impossible to conduct several at once." (Kahneman, 2011, p. 23)

The most obvious way to do this is by practice. Another way to free up resources is to minimize the self-control needed to keep your attention on what you're doing, by being in a state of 'flow', as described in the previous section. Flow, which is part of developing automaticity, reflects System 1 thinking. But such abilities cannot be achieved without a fair amount of help from System 2 thinking, to which we now turn.

Probably one of the most commonly understood aspects of skill acquisition is that acquiring a skill takes "practice, practice, practice". How much practice? Frequent estimates place the amount of time necessary to achieve expertise in any field at 10 years or 10,000 hours.<sup>24</sup> However, mere experience isn't sufficient for achieving expertise. People reach a certain level of acceptable performance, after which further experience does not lead to any improvement in performance. Additional experience may make performing at that level of skillfulness easier, but that is not the same as actually improving one's performance. Thus, the number of years of experience one has is not a sufficient predictor of performance. While having 10 years of experience may be necessary for expertise, it does not by itself guarantee expertise.

What more is needed? Research indicates that a particular kind of experience is necessary for expertise, as it turns out that the quality of the practice matters just as much as the quantity. Improving your level of skill requires not the mere repetition of things you already know how to do, but continually striving to do things that you currently cannot do. This kind of experience is referred to as 'deliberate practice', and it's roughly 10,000 hours of deliberate practice that's needed for expertise. Deliberate practice requires having specific goals in mind for improvement, rather than a more general goal of 'getting better'. There need to be specific aspects of your performance that you go about planning how to improve, which then structures the kind of deliberate practice you engage in<sup>25</sup>. As you engage in deliberate practice you seek out feedback about your performance, in the hopes of identifying and correcting errors. You keep monitoring your progress as you practice. If you don't seem to be progressing, you may need to redesign your practice sessions. If instead you keep up a steady progression, then at some point you reach your goal. At that point it's time to set out a new goal to strive to accomplish. This is how you improve upon your current level of performance.

<sup>24</sup> Horn and Masunaga, 2006, p. 601

<sup>25</sup> Horn and Masunaga, 2006, p. 601

Deliberate practice clearly involves System 2 thinking, as “the requirement for *concentration* sets deliberate practice apart from both mindless, routine performance and playful engagement, as the latter two types of activities would, if anything, merely strengthen the current mediating cognitive mechanisms, rather than modify them to allow increases in the level of performance.”<sup>26</sup> This can also be seen in the fact that in addition to getting feedback from others<sup>27</sup>, you need to learn how to monitor your own performance while practicing.

Self-regulation is important in acquiring expertise because feedback cannot come merely from others, as crucial as that is in the early stages of skill acquisition. “Because high levels of skill must be practiced and adapted personally to dynamic contexts, aspiring experts need to develop a self-disciplined approach to learning and practice to gain consistency.”<sup>28</sup> Often there won’t be a coach around when you are exercising your skill, and so you need to learn how to provide yourself feedback on your performance. Therefore, it is important for deliberate practice that you are able to monitor your own behavior during such sessions, so that you can provide feedback for yourself.

Experts need to not only monitor their own behavior, but they also must monitor the environment that they are working in for changes.<sup>29</sup> This is especially relevant when experts face situations that contain features they have little prior experience with. Because expertise develops out of concrete experience, experts will be at their best when facing relatively familiar situations. Thus, experts also need to be aware of when they are facing situations that include unique features, so as to adjust their performance. While they may not perform as well in truly unique situations, they will still fare better than novices. There is a bit of a balancing act that has to be performed between automaticity and monitoring one’s environment, as experts still need to be aware of their situation in order to detect features that may be out of the

<sup>26</sup> Ericsson, 2006c p. 692

<sup>27</sup> This opens up the possibility that one makes a correct choice, but the feedback from other people is that the choice was wrong (and this may be simply a mistake or an intentional attempt to discredit the choice).

<sup>28</sup> Zimmerman, 2006, p. 706

<sup>29</sup> “This kind of monitoring prevents blind alleys, errors, and the need for extensive back-up and retraction, thus ensuring overall progress to a goal. In addition, these same kinds of monitoring behaviors are critical throughout the process of acquiring knowledge and skills on which expertise depends.” (Feltovich, Prietula and Ericsson, 2006, p. 56)

ordinary. But even this kind of situational awareness can go on with little conscious effort.<sup>30</sup>

Furthermore, once expertise has been achieved in a skill, the same kind of deliberate practice and self-monitoring is necessary to retain expert performance. While everyday wisdom teaches us that once you achieve expertise you never really lose it, research on age and expertise shows that “maintaining skills is as effortful as acquiring them in the first place”.<sup>31</sup> Expertise requires some level of routine practice to maintain it or the level of skill degrades over time.

Given the overall difficulty of achieving expertise, one of the most important factors for determining whether someone can attain that level of performance is motivation. “Unless a person wants to pursue the difficult path that leads to the development of talent, neither innate potential nor all the knowledge in the world will suffice.”<sup>32</sup> Not only does an aspiring expert need to be strongly motivated to perform well, in the face of adversity and over a long period of time, but even after achieving expertise a high level of motivation is still required to maintain one’s expertise.

Having the requisite motivation is not purely an individual affair though. A variety of social factors come into play in supporting expertise: “Becoming an expert in almost anything requires literally years of work. People will do this only if they have some initial success, enjoy the work, and are supported by the social climate. Expertise is not solely a cognitive affair” (Hunt, 2006, p. 36).

Social support can either help or hinder the development of expertise in certain domains (say by devoting public funds towards education and training in certain fields), as well as an individual’s motivation to achieve expertise. One key group in the social support of expertise is parents. Since the development of expertise takes considerable time it helps to start young, and supportive parents can have a big impact on this stage of development. Teachers are another key group that can help to motivate students in acquiring skills. Not only is encouragement helpful, but also one’s social class in terms of affording coaches and training.<sup>33</sup> Social support can of course work both ways, as there

<sup>30</sup> “Evidence exists, for example, that metacognition can be automatic (Reder & Shunn, 1996), thus avoiding Tulving’s (1994) consciousness requirement for metacognitive judgement.” Feltovich, Prietula and Ericsson, 2006, p. 57)

<sup>31</sup> Krampe, 2006, p. 733

<sup>32</sup> Csikszentmihalyi, Rathunde and Whalen, 1993, pp. 31-32

<sup>33</sup> “A child’s acquisition of expertise in both common and more esoteric activities emerges from modeling, instruction, monitoring, and guidance activities by his or her parents, teachers, and peers within the social

are ways in which society can discourage individuals and groups from the achievement of expertise. It is important to note that there are several issues related to social support for expertise, such as social biases (such as gender, race, and social class) in who is publicly recognized as an expert; and in respect to supporting groups and individuals in their striving to attain expertise.<sup>34</sup>

### 3. Dreyfus on Expertise

The Dreyfus model of skill acquisition is the most well-known philosophical account of expertise. It represents a phenomenological approach to understanding expertise. While the psychological research in general supports their view, it also provides a more balanced picture of expertise than you otherwise find on the Dreyfus model, which emphasizes the System 1 (intuitive) aspects of expertise while neglecting the System 2 (deliberative) aspects. The following section presents a very brief outline of the Dreyfus model, and then points out those features of expertise that the model neglects.

The Dreyfus model divides skill acquisition into five stages: novice, advanced beginner, competent performer, proficient performer, and expert. At the initial stages of skill acquisition, novices follow simple and context-free rules, such as, in cases of driving, “shift into second gear at ten m.p.h.,” or use the two-second rule in judging how much space to leave between you and the car in front of you. Since the rules at this stage are context-free, however, they are apt to fail in a variety of different circumstances, such as when driving in the rain or in heavy traffic. As the novice gains experience, she discovers new features of situations, or someone else points them out, as relevant. Instead of relying only upon rules, the advanced beginner starts using maxims, which are not context-free like rules, but rather take into account the new features of situations of which the advanced beginner is aware. A maxim for driving might be “when the engine sounds like its racing shift up in gear.” This maxim refers to the situational aspect of engine sounds, which it takes experience to recognize, and so this type of instruction is inappropriate for novices.

milieu of the family, the school, and the community. . . Bloom (1985) found that their parents not only nurtured the child’s initial interest and provided or arranged high-quality instruction, they also emphasized the importance of dedicated practice: “To excel, to *do one’s best*, to *work hard*, and to *spend one’s time constructively* were emphasized over and over again” (p. 10).” (Zimmerman, 2006, p. 706)

<sup>34</sup> Further issues regarding the social dimensions of expert status are taken up in the last section.

Even these maxims have their limitations, however, for the number of situational factors can become overwhelming. Moving beyond maxims requires making choices about what the most relevant factor is in a situation, and this is done by adopting a specific plan or perspective. According to the Dreyfus model, the competent performer feels responsible for both the choice of perspective and the outcome of that choice, and thus becomes emotionally involved in the experience of the outcome. “An outcome that is clearly successful is deeply satisfying and leaves a vivid memory of the plan chosen and of the situation as seen from the perspective of the plan. Disasters, likewise, are not easily forgotten.”<sup>35</sup> These outcomes provide the feedback that a person needs in order to improve her skill. The feedback, if positive, reinforces making that choice again in a similar situation. The feedback, if negative, prompts the person to make a different choice in that situation.

While the competent performer has to make up rules to help her decide what plan or perspective to adopt in order to focus in on the relevant features of a situation, the proficient performer no longer uses rules or even makes a choice about a plan. The proficient performer simply experiences the situation in the light of a certain perspective, without making a conscious decision about the most appropriate perspective to take in the situation.

The final stage is that of expertise. Dreyfus discovered that one of the hallmark features of expertise is an intuitive form of decision-making. By ‘intuition’, he is “referring to the understanding that effortlessly occurs upon seeing similarities with previous experiences.”<sup>36</sup> The ability of the expert to act well intuitively is due to the expert’s experience and familiarity with the situation in which she acts. The immediacy of the expert’s judgment occurs because of repeated exposure to similar previous experiences, and the outcome of actions taken in those situations, so that:

With enough experience with a variety of situations, all seen from the same perspective but requiring different tactical decisions, the proficient performer seems gradually to decompose this class of situations into subclasses, each of which share the same decision, single action, or tactic. This allows an immediate intuitive response to each situation. (Dreyfus and Dreyfus, 1991, p. 235).

<sup>35</sup> Dreyfus and Dreyfus 1986, p. 26

<sup>36</sup> Dreyfus and Dreyfus, 1986, p. 28

The expert knows what actions are required and how to perform them in that situation without detached calculation or having to weigh alternatives. An expert driver will shift gears when appropriate without even being aware of it. Dreyfus also found that experts frequently were not able to give an account of how they knew what to do. One might be an expert skier, but find it quite difficult to teach others how to ski. On the Dreyfus account, since experts generally act well without applying rules and principles, it is no surprise that experts often find it difficult to explain their actions by reference to principles. Of course some experts are articulate or are good at teaching others, but these abilities are not in any way necessary for expertise.

The Dreyfus account is at its best when discussing expert performance as involving automaticity and intuitive judgment, based on experience and pattern recognition. However, Dreyfus seems sometimes at pains to avoid talk of deliberation and choice with regard to expertise. It's mainly at the early stages of skill acquisition that Dreyfus acknowledges the role of deliberation and the need to make conscious choices between alternatives (i.e. characteristics of System 2 thinking). Dreyfus is correct to note that there's a definite change in performance past the stage of competency, where the higher levels of performance are characterized more by automaticity and System 1 processing. However, Dreyfus underestimates the degree of System 2 thinking in skill acquisition in two respects: 1) the degree to which practice must be 'deliberate practice' involving self-regulating behavior in the early stages of skill acquisition; and 2) that deliberate practice and self-regulating behavior carries over into advanced stages of skill acquisition (even though rule-following does not).

First, the initial stages of skill development on the Dreyfus model are characterized in terms of rule-following. A novice relies on context-free rules, at least until she gains enough experience that she can use more sophisticated rules that refer to situational cues that she has learned to recognize. While this is certainly a familiar aspect of learning a skill, what isn't mentioned is the role of deliberate practice and the need for self-regulating behavior. That is, the focus on the Dreyfus model is what performance is like at each stage of skill development, rather than the factors that enable one to improve (beyond needing more experience). Second, the main deliberative factor in expertise on the Dreyfus model is the rule-following of the novice and advanced beginner, which is something that needs to be left behind to progress to higher levels of skill. Thus, on the Dreyfus model, there seems to be no important deliberative

aspects that carry over into higher levels of performance. However, improvement in one's level of skill always requires deliberate aspects such as deliberate practice and self-regulating behavior. So while rule-following does drop out of the picture at higher levels of skill, not all deliberative aspects from the initial stages drop out. These are important features of expertise that are absent in the Dreyfus model, and certainly skew their view towards an overemphasis on System 1 aspects to skill acquisition and expert performance.

#### 4. Annas on Expertise

Julia Annas frequently discusses practical skills, as she works in virtue theory and sees many structural similarities between the acquisition of virtue and that of practical skills. What is of concern here is the view of skills and expertise that she's working from, rather than her account of virtue.<sup>37</sup> She relies on an intellectual account of expertise, which portrays expertise as more of a matter of System 2 thinking. Her view can thus help correct for how the Dreyfus account underemphasizes the role of deliberative processes in expertise. Though, like Dreyfus, at times she overstates her case, and a few of her claims are not well-supported by the current psychological research. Annas's discussion draws mainly on Socratic ideas about the nature of skills and expertise. According to Annas, there are three necessary elements of a genuine skill: the skill must be teachable, there must be unifying principles underlying the skill that the expert can grasp, and that experts can give an account of skilled actions.

The first element is that the skill is teachable. Since the expert has learned something, she should be able to teach what she has learned to someone else. The expert has learned the theory behind the skill. This contrasts with what Socrates refers to as a 'knack', which is something that can be picked up merely by trying to do it yourself, or by watching someone else do it. Knacks lack the intellectual component that is found in skills. Rhetoric and cooking are putative examples of mere knacks. Genuine skills have a strong intellectual component, and this is what the expert is able to teach.

Annas is surely right that skills are teachable, as coaches and trainers can provide essential feedback and deliberate practice routines for improvement. However, Annas is also making a stronger claim, that the expert, in virtue of

<sup>37</sup> For an example of applying expertise to virtue, see Stichter, 2011.



her expertise, should be able to teach others. The psychological research does not support this stronger claim:

Although it is tempting to believe that upon knowing how the expert does something, one might be able to “teach” this to novices directly, this has not been the case (e.g. Klein & Hoffman, 1993). Expertise is a long-term developmental process, resulting from rich instrumental experiences in the world and extensive practice. These cannot simply be handed to someone. (Feltovich, Prietula and Ericsson, 2006, p. 46).

Part of the difficulty in teaching is that expertise is not primarily an intellectual grasp of theory, but the development of a number of cognitive adaptations that result from experience and practice. There’s no way to gain this kind of knowledge except by going through the same kind of process. Furthermore, one weakness of experts appears to be that they have trouble predicting novice performance, perhaps because they cannot easily take on the perspective of a novice attempting a task. Certainly a good teacher needs to be able to appreciate the perspective of a novice, in order to provide helpful guidance at that stage of skill development. Being an expert seems to carry an inherent disadvantage in that regard. However, one should not overstate the case, as experts can be good teachers. It’s just not the case that expertise translates necessarily into being able to teach well. For example, many in academia should be familiar with professors that are good at research, but not so good at teaching.

The second element expands upon the intellectual component found in teaching. To possess a skill requires what Annas refers to as having “a unified grasp of its field.” (Annas, 1995, p. 231). This implies that there are principles that unify the field of a skill, and that the expert has a grasp of these principles. There is no such thing as having expert knowledge of only part of the field. One could not claim to be an expert at something as narrow as only being able to fix Toyotas, or to claim, as Ion does, only to know Homer and not much of any other poet. Annas recognizes that:

This probably surprises us. Surely, we think, Ion does have expert knowledge, only not enough: he knows only part of the field. But Socrates does not accept this way of looking at it. If there is such a thing as the skill that consists in mastering poetry, then it consists in grasping the principles which apply over the whole field. To fail to do this in one area reveals that one cannot do it at all. [...] In each case the skill in question is one that you do not have until you have mastered all the relevant elements in the field. (Annas, 1995, pp. 231–232).

Expertise requires understanding the principles that govern the entire field, and not just some parts of it. This unified grasp is what allows experts to deal with unfamiliar situations in the way that someone who has simply memorized a set of rules cannot, since it enables them to act well with regard to all areas of the field.

However, the expertise literature throws some doubt on experts having a ‘unified grasp of the field’. It’s not just the case that expertise is domain-limited, but it’s also limited even within a domain. Although Annas relies on medicine as the main example of a ‘genuine’ skill that fits her account of skills:

studies showed that the same physician can demonstrate widely different profiles of competence, depending on his or her particular experiential history with different types of cases. Indeed, in modern medical education, where assessment of clinical skill is often evaluated by performance on real or simulated cases, it has been found that because of the case-specificity of clinical skill, a large number of cases (on the order of fourteen to eighteen) are needed to achieve an acceptably reliable assessment of skill. (Feltovich, Pricetula and Ericsson, 2006, p. 47).

Expertise arises out of experience, and one’s experience places a limiting factor on which situations one can display expert-level performance. Psychological research emphasizes that experts rely on contextual cues, such that when the expert is operating in an unusual context, they lack information necessary to perform at an expert level.

For example, in a medical domain, experts seem to rely on the tacit enabling conditions of a situation for diagnosis (Feltovich & Barrows, 1984). The enabling conditions are background information such as age, sex, previous diseases, occupation, drug use, and so forth. These circumstances are not necessarily causally related to diseases, but physicians pick up and use such correlational knowledge from clinical practice. [...] The implication is that without the contextual enabling information, expert physicians might be more limited in their ability to make an accurate diagnosis (Chi, 2006, p. 25).

Annas is correct in thinking that experts should be able to generalize to some extent from their experiences, such that they would have some idea of how to act well in unfamiliar situations. But it is important not to overestimate how well experts will react in novel situations, since their expertise is still linked to a

certain history of experience.<sup>38</sup>

The third element of a genuine skill further develops the previous intellectual components, by requiring that experts have the ability to ‘give an account’ of their actions. Giving an account, according to Annas means “that the person with a skill be able explicitly to explain and justify her particular decisions and judgements, and to do so in terms of some general grasp of the principles which define that skill.” (Annas, 1995, p. 233). The expert needs to be able to articulate the reasons for her actions, and this explanation should draw upon the expert’s grasp of the principles underlying the skill. Although this condition could be thought of as requiring merely that the principles are articulatable, rather than requiring that the expert can actually articulate the reasons herself, Annas explicitly describes this requirement in terms of the expert being able to articulate the reasons for her actions.

As discussed earlier, research shows that experts cannot be relied on to accurately articulate their reasons for action. However, in her defense, Annas is aware that the three essential elements form a high intellectual standard for skills that strikes people as counterintuitive. She notes that:

This idea, that conveying and acquiring a skill requires articulacy, often meets resistance. This may take the form of pointing to skills where articulacy does not appear to be necessary; sometimes gardening is given as an example. In some cases, such as physical skills, the person outstanding in the skill may not be the best at conveying it (as with athletes and coaches). Many of these will be cases where what is at stake is really mastery of technical matters needed for the exercise of the skill, or where what is important is natural talent. (Annas, 2011, p. 19).

Annas admits that the requirement of giving an account is not true for a number of actual skills. For Annas, this result is not problematic, so long as there are some skills that do display these strong intellectual components, because it’s those kinds of skills that she thinks share a structure similar to virtue. So her claims should not be read as applying to everything we might label a skill.

<sup>38</sup> “For example: “The experienced pilot who has never encountered or been trained for a particular anomaly will be challenged to process information in working memory to determine what is happening, and may be inefficient in searching for relevant information to solve the problem, in much the same way as when she was a novice pilot (although it is likely that she will not be as bad off as a complete novice). Most people do not operate at the level of novice all the time or expert all the time, but rather move around in between, using combinations of cognitive mechanisms depending on the situation at hand and the availability of key constructs (e.g. mental models and schema).” (Endsley, 2006, p. 640)

That there are skills that display the three intellectual components, however, is far from clear. One of the skills she does mention is medicine, but evidence from the medical field calls into question experts displaying all of these strong intellectual components. For example:

Bias is probably one of the most serious handicaps of experts, especially in the medical profession. . . . my colleagues and I found the experienced physicians to manifest serious biases. We presented several types of cases to specialists, such as hematologists, cardiologists, and infectious disease specialists. Some were hematology cases and others were cardiology cases. We found that regardless of the type of specialized case, specialists tended to generate hypotheses that corresponded to their field of expertise: Cardiologists tended to generate more cardiology-type hypotheses, whether the case was one of a blood disease or an infectious disease (Hashem, Chi, & Friedman, 2003). This tendency to generate diagnoses about which they have more knowledge clearly can cause greater errors. Moreover, experts seem to be more susceptible to suggestions that can bias their choices than novices (Walther, Fiedler, & Nickel, 2003). (Chi, 2006, pp. 26-27).

The psychological research thus appears to temper some of Annas's claims about expertise, when she overstates the intellectual aspects.

Finally, there is one other reason to be concerned about overemphasizing the role of articulation in expertise. Patricia Benner<sup>39</sup> carried out studies of experts in the field of nursing. When studying nurses with a track record of life-saving decisions in emergency situations, she found that often the nurses could not fully articulate how they knew what to do. Benner quotes an expert psychiatric nurse clinician who is talking about her clinical judgments:

When I say to a doctor, “the patient is psychotic,” I don't always know how to legitimize that statement. But I am never wrong. Because I know psychosis from inside out. And I feel that, and I know it, and I trust it . . . One of the things that I am doing now is getting some in-service in to talk to us about language. But all I am really trying to do is find words within the jargon to talk about something that I don't think is particularly describable. (Benner, 2001, p. 32).

If we view articulation as a necessary component of expertise, then this nurse would seemingly not count as an expert because she is not able to give an articulate justification for her clinical judgment. But this would reach the

<sup>39</sup> Benner was applying the Dreyfus model in particular.

wrong conclusion, as Benner informs us that this is a nurse who has over 15 years' experience in the field and who is reliably correct. One of the most serious problems for the nurses is that their judgments were not taken as seriously as doctors because of an assumption that their lack of articulation signaled a lack of knowledge, and so they were also accorded less power and status within the hospital. So it matters a great deal that we get an accurate picture of what really goes into acquiring skills and expertise. There are important intuitive and deliberative aspects to both skill acquisition and expert performance. The psychological research helps to correct those philosophical accounts of expertise that overemphasize one aspect over the other.

### 5. Expertise *versus* Experts

In this final section, the discussion transitions from expertise to experts. It might be thought that an expert is just someone who has expertise, but it's important to distinguish between having expertise in a field and being credited as an expert. As discussed above, expertise refers specifically to the highest levels of skill that human beings have reached. Acquiring a skill is a matter of getting better at some task, and so the possession of skill is a matter of degree. Progression through the stages of skill acquisition is mostly a matter of experience and practice. While there are frequently social aspects to becoming skilled, such as who is encouraged to pursue different skill sets or who can afford access to training, acquiring expertise in a domain does not require being credited as an expert. For example, one could be a complete novice at fishing, be stranded alone on a desert island for 10 years, and with practice and experience during that time, develop a high level of skill at fishing, and even achieve expertise in it. So expertise, as I refer to it here and as it's studied in the psychological literature, is at its core a study of the phenomenon of how people get better at tasks.

Being credited as an expert is not the same thing as having expertise. There are a different set of questions that get raised when inquiring about who should be credited as an expert. Presumably the point of crediting people as experts is typically that there is something that we want from them. We hope that those we deem to be experts accurately tracks those who have achieved expertise in a field, but the two can come apart (perhaps due to bias or incomplete information). You can have people credited with expert status that lack expertise, and people with expertise may not be granted expert status. One of the main concerns with experts is that being credited with expert status can confer power and authority

that is distinct from the expertise itself. So while it's relatively uncontroversial that people can achieve expertise, what is controversial is our attempt credit people with expert status.<sup>40</sup> With respect to determining who to confer expert status upon, there are a host of social, political, and epistemological concerns.

While addressing all these concerns is far beyond the scope of this paper, in doing so it will be important to keep in mind not only what expertise is like, but also the difference between expertise and being credited as an expert. Deciding who to credit as an expert is not like describing a natural kind. We credit people with expert status because it serves a useful function, and we decide who to confer this status upon depending on what we want from them, and this also determines what power we choose to grant to them. So we want to avoid what some have termed an “immaculate” conception of the expert – someone who counts as an expert whether we like it or not, and to whom we must defer judgment.<sup>41</sup> While someone can have expertise without any social validation, the same is not true with regard to possessing expert status.

Not keeping this distinction in mind can lead to a very misleading picture of what it is to be an expert. No doubt some putative expert might want to push the immaculate conception on us, but we should resist it. While someone may be able to demonstrate their expertise, it doesn't follow from that demonstration that we should accord that person the status of an expert. If someone loudly declares “I'm an expert”, then we can always reply “Only if we say you are”. Expert advice can be ignored, and expert status can be revoked, since ultimately we're doing this (if at all) because we find it useful. In addition, seeing a few of the limitations of expertise, as mentioned above, can be helpful in reminding us of the fallibility of those who have achieved expertise. One hope for this account is that it not only helps to inform discussions of experts, but also to contest and challenge expert discourse.

<sup>40</sup> For these reasons, I try to talk in terms of crediting people with expert status, to emphasize that we confer this status on others (along with whatever power goes along with it). I try to avoid talking in terms of ‘recognizing’ or ‘identifying’ experts, as that may suggest people are experts independent of us conferring that status on them.

<sup>41</sup> See Turner, 2001.

## REFERENCES

- Annas, J. (2011). *Intelligent Virtue*. Oxford: Oxford University Press.
- Annas, J. (1995). Virtue as a Skill. *International Journal of Philosophical Studies*, 3 (2), pp. 227- 243.
- Benner, P. (2001). *From Novice to Expert*. New Jersey: Prentice Hall Health
- Chi, M. (2006). Two Approaches to the Study of Experts' Characteristics. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 21-30). Cambridge: Cambridge University Press.
- Csikszentmihalyi, M. and Rathunde, K. and Whalen, S. (1993). *Talented Teenagers: The roots of success and failure*. New York: Cambridge University Press.
- Dreyfus, H. and Dreyfus, S. (1991). Towards a Phenomenology of Ethical Expertise. *Human Studies*, 14, pp. 229-250.
- Dreyfus, H. & Dreyfus, S. (1986). *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Oxford: Basil Blackwell.
- Endsley, M. (2006). Expertise and Situation Awareness. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 633-651). Cambridge: Cambridge University Press.
- Ericsson, K.A. (2006a). An Introduction to *Cambridge Handbook of Expertise and Expert Performance*: Its Development, Organization, and Content. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 3-20). Cambridge: Cambridge University Press.
- Ericsson, K.A. (2006b). Protocol Analysis and Expert Thought: Concurrent Verbalizations of Thinking during Experts' Performance on Representative Tasks. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 223-241). Cambridge: Cambridge University Press.
- Ericsson, K.A. (2006c). The Influence of Experience and Deliberate Practice on the Development of Superior Expert Performance. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 683-704). Cambridge: Cambridge University Press.
- Evans, J. and Stanovich, K. (2013). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science*, 8 (3), pp. 223-241.
- Feltovich, P, Prietula, M. and Ericsson, K.A. (2006). Studies of Expertise from Psychological Perspectives. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 41-68). Cambridge: Cambridge University Press.

- Horn, J, Masunaga, H. (2006). A Merging Theory of Expertise and Intelligence. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 587-612). Cambridge: Cambridge University Press.
- Hunt, E. (2006). Expertise, Talent, and Social Encouragement. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 31-38). Cambridge: Cambridge University Press.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kahneman, D. and Klein, G. (2009). Conditions for Intuitive Expertise: A Failure to Disagree. *American Psychologist*, 64 (6), pp. 515-526.
- Keren, G, and Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on psychological science*, 4, pp. 533-550.
- Klein, G, Ross, K. and Shafter, J. (2006). Professional Judgments and “Naturalistic Decision Making”. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 403-419). Cambridge: Cambridge University Press.
- Krampe, R. and Charness, N. (2006). Aging and Expertise. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 723-742). Cambridge: Cambridge University Press.
- Lapsley, D. and Hill, P. (2008). On Dual Processing and Heuristic Approaches to Moral Cognition. *Journal of Moral Education*, 37 (3), pp. 313-332.
- Lapsley, D. and Narvaez, D. (2005). The Psychological Foundations of Everyday Morality and Moral Expertise. In D. Lapsley and C. Power (Eds.), *Character Psychology and Character Education* (pp. 140-165). Notre Dame: IN: University of Notre Dame Press.
- Stichter, M. (2011). Virtues, Skills, and Right Action, *Ethical Theory and Moral Practice*, 14, pp. 73-86.
- Turner, S. (2001) What is the Problem with Experts? *Social Studies of Science*, 31 (1), pp. 123-149.
- Zimmerman, B. (2006). Development and Adaptation of Expertise: The Role of Self-Regulatory Processes and Beliefs. In K. A. Ericsson (Ed.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 705-722). Cambridge: Cambridge University Press.