6-2013

A Cognitive Approach to Teaching Strategies

Emily Esch
College of Saint Benedict/Saint John's University, eesch@csbsju.edu

A Cognitive Approach to Teaching Philosophy

EMILY ESCH

College of St. Benedict, St. John's University

Abstract: Our knowledge of how the mind works is growing rapidly. One area of particular interest to philosophy teachers is research on reasoning and decision making processes. I explore one model of human cognition that offers new ways of thinking about how to teach philosophical skills. The bulk of the paper is dedicated to exposition of the model and the evidence that supports it; at the end of the paper, I suggest ways these findings might be incorporated into the classroom.

I. Introduction

In the 1970s, the psychologists Amos Tversky and Daniel Kahneman began an influential research program into the intuitions that people use to form judgments and make decisions, intuitions that became known as heuristics and biases. Heuristics are cognitive strategies that people use to make judgments; for example, people tend to use whatever information is most easily available to them. Here is an illustration of this tendency, which is now known as the availability heuristic (Kahneman 2011: 7). When asked if there are more words that begin with a "k" or more words that have "k" as the third letter, the majority of people respond incorrectly that more words begin with a "k." The strategy people are using is to come up with lists of words with "k" as the first and "k" as the third letter. If you try it, you'll find is easier to come up with words that begin with "k" (or any other letter) than words that have "k" as the third letter. Hence, the availability heuristic leads people to erroneously conclude that there must be more words that begin with "k." Sometimes these cognitive strategies work well, but other times they lead to biases. Relying on the availability heuristic, for example, will often lead to discounting statistical evidence. The availability of stories of children being kidnapped by strangers leads many parents to overestimate the likelihood of their own children being kidnapped.

Building on the work of Tversky and Kahneman, researchers have found that humans make systematic mistakes on a number of cognitive tasks. The psychologist Keith Stanovich describes some of the systematic errors as follows: people "assess probabilities incorrectly, they test hypotheses inefficiently, they violate the axioms of utility theory, they do not properly calibrate degrees of belief, their choices are affected by irrelevant context, they ignore the alternative hypothesis when evaluating data" (Stanovich 2011: 7). Stanovich (along with collaborators Richard West and Maggie Toplak) has developed a taxonomy of these errors and uses this taxonomy to build a model of human cognition (Stanovich 2011, Stanovich 2009). But, as Stanovich points out, there are people who don't make these kinds of systematic errors. If we look at what these individuals are doing differently, we might be able to understand how to improve the average person's reasoning abilities. It is in this context that Stanovich introduces his model of cognition.

II. Dual Processing and the Tripartite Model

One of the most fruitful research paradigms of the last several decades is the dual processing model (Stanovich 2011: 16–25, Kahneman 2011: 19–108; Hassin, Uleman, and Bargh 2006). The dual processing model proposes that cognition can be divided into two types of systems, *automatic* and *effortful*.1 Both of these types are familiar. We have all had the experience of behaving on autopilot: finding ourselves in our office without recalling the walk across campus, recognizing with a start that miles have passed unknowingly on a long distance drive, or playing a familiar game of catch while maintaining a conversation. In each of these cases, walking, driving, and catching, the activity can be accomplished without paying it much attention. The automatic system is also responsible for our ability to read anger on a human face, our thoughts of "2" when shown "1+1," and our feeling of compassion when shown distressful photos. Automatic processing is automatic, fast, and not dependent on higher order processing. Several automatic processes can be run simultaneously. What scientists have discovered over the past few decades is that the automatic system seems to explain our behavior much more often than most of us would have expected.

But, of course, we are not always on autopilot. Much of our cognitive life consists in consciously thinking about the decisions in front of us, like deciding where to go on vacation, planning our daily schedule, or answering the problem "378+294." This kind of conscious deliberating is slow, serial, and cognitively demanding. Effortful processes are cognitively expensive. They tend to run serially, not in parallel. So multitasking when effortful processes are engaged leads to a decline in function.

Tversky and Kahneman explore in detail the way the automatic system makes judgments and decisions. We'll see examples of their findings below. But first let me make a point about the big picture. As Stanovich points out, with Tversky and Kahneman's work in mind, one of the dominant themes running through the last four decades of psychological research is the notion that humans are "cognitive misers" (Stanovich 2011: 29). Kahneman, himself, says, we "conduct our mental lives by the law of least effort" (Kahneman 2011: 38). The automatic processes require less computational capacity, so there is more power left over to engage in other tasks. For what seem to be adaptive reasons, humans default to automatic processing whenever they have a chance.2 In other words, and not at all surprising to teachers of philosophy, given a chance to be intellectually lazy, human beings take it. Overcoming this tendency takes cognitive effort. Pedagogically, this seems an important point to emphasize. It is not students who are cognitively lazy; it is humans.

Figure 1. Illustration of Stanovich's tripartite model.

Stanovich's contribution to the dual processing model is to complicate our understanding of the effortful processes, illustrated in figure 1. Stanovich adds to the dual processing model by splitting apart the effortful processes into *algorithmic processing* and *reflective processing*, turning the dual processing model into a tripartite model (Stanovich 2011: 29–46). This distinction is best understood by looking at various types of testing, discussed below, but here is a quick conceptual gloss. Algorithmic processes sustain our working memory and our ability to form and manipulate representations "offline," that is, unconnected to our behavior. Underlying the algorithmic processes is the ability to decouple our representations from the world. The ability to decouple our representations from our behavior and from what we currently perceive allows us to imagine possibilities that are not actual, e.g., a philosophical thought experiment or a future plan of action. As many have noted, this kind of hypothetical reasoning is a central part of human cognition.

Reflective processes are of a higher order than algorithmic processes: they involve a person's epistemic values and her ability to self-regulate (Stanovich 2011: 36). Whereas the algorithmic processes embody the nuts and bolts of the operation, the reflective processes determine how the machine is used. Algorithmic processes are measured by skills, while the reflective processes are measured by thinking dispositions. These dispositions include

> the tendency to seek various points of view before coming to a conclusion, the disposition to think extensively about a problem before responding, the tendency to calibrate the degree of strength of one's opinion to the degree of evidence available, the tendency to think about future consequences before taking action, the tendency to explicitly weigh pluses and minuses of situations before making a decision, and the tendency to seek nuance and avoid absolutism. (Stanovich 2011: 36)

To state the obvious, many of these are thinking dispositions that we try to inculcate in our students. For anyone who knows the literature on critical thinking, this list will be familiar. The same kinds of dispositions are found in the description of the ideal critical thinker composed by the Delphi project, a committee of dozens of researchers sponsored by the American Philosophical Association (Facione 2000).

As I mentioned, the distinction between algorithmic and reflective processing is best understood by examining differences in the kinds of tests designed to measure cognition. I'll begin with a brief description of the most familiar of intelligence tests, the IQ test. IQ tests (and proxies like the SAT) are designed to measure general intelligence, called *G*. Two dominant factors reflected in the generalized score are *fluid intelligence* (Gf) and *crystallized intelligence* (Gc) (Stanovich 2011: 52–53; Flynn 2007). Gf measures reasoning abilities on novel tasks, like analogies and nonverbal pattern finding. Gc measures declarative knowledge, like vocabulary and general knowledge.[3] Following the practice of the psychological literature, I will use the term "intelligence" to mean whatever it is that is measured by IQ tests.[4]

Significantly, IQ tests do not attempt to measure any of the reflective processes (Stanovich 2011: 96–97). Not only are the two processes conceptually distinct, but they are also psychometrically distinct. Stanovich claims that the difference between algorithmic and reflective processing is reflected in the difference between *maximal (or optimal) performance situations* and *typical performance situations*. IQ tests are conducted under maximal performance situations. In maximal performance situations, the participant is told what the goal of the task is and how to best maximize her performance (Stanovich 2009: 31). These are usually timed tests designed to measure efficiency. Typical performance situations more closely mimic situations one might encounter in the real world. The tasks assigned in these situations are more open ended: no goals are assigned and no specific instructions on how to best complete the task are given. These situations are designed to measure motivation and dispositions. Critical thinking tests are conducted under typical performance situations and they assess thinking dispositions (Stanovich 2011: 40).

Figure 2. A diagram illustrating the relationships between maximal and typical performance situations and testing.

Examining some specific examples will give us a deeper understanding of the division between typical and maximal conditions, which in turn elicit different kinds of processes in the thinker:

According to Stanovich, when syllogistic reasoning shows up on an IQ test, it looks like this:

1. All As are Bs.
2. C is B.
3. Therefore, C is A.

The argument is presented in terms of As and Bs in order to remove the possibility of someone's prior beliefs influencing her answer. IQ test designers try to eliminate the need for reflective level thinking as much as possible. They are trying to isolate and measure only the algorithmic processes, though this is difficult to do and they are not always successful (Stanovich 2011: 41).

Here is the kind of item one might find on a critical thinking test, the type of test that measures reflective processes (Stanovich 2011: 40). Without being given explicit instructions, participants are asked whether the following is a valid argument:

1. All living things need water.
2. Roses need water.
3. Therefore, roses are living things.

This kind of example tests reflective processes, since it pits prior beliefs (roses are living things) against the evidence presented in the argument. Participants who default to automatic reasoning, i.e., recognize the conclusion is true so accept the argument without reflection, will get it wrong. About 70 percent of university students given this problem think that the argument is valid (Stanovich 2011: 40).

We are now at the point where I can explain why I find Stanovich's research so interesting for teachers. It turns out, surprisingly, that on many (though not all) of the critical thinking tasks that have been studied there is little to no correlation between completing the task correctly and intelligence as measured by IQ.5 In other words, efficiency in algorithmic processing does not always explain high scores on a critical reasoning test (Stanovich 2011: 140). Even on those tasks where there is a correlation between IQ and the critical thinking task, "the magnitude of the associations with [IQ] left considerable room for the possibility that the remaining reliable variance might index systematic variation in reflective-level psychological tendencies" (Stanovich 2011: 44–45). Cognitive ability as measured by IQ is, at best, only part of the story behind successful reasoning.6

The other part of the story, reflective processing, Stanovich calls *rationality.* The kinds of epistemic dispositions associated with reflective processing—slowing down and thinking through all sides of a problem, balancing the level of conviction with the strength of the evidence, appreciating nuance and avoiding dogmatism—are all tendencies that can be taught, and, in fact, are currently taught in most philosophy courses. But, despite lots of first-hand experience with smart people being stupid, until reading Stanovich I had been working under an unexamined assumption that there was a close correlation between algorithmic processes and reflective processes. Learning that these are, by and large, separate kinds of processes has made me more optimistic about my students' ability to make real progress in developing reasoning skills in a short period of time (e.g., a semester). It has also made me more sympathetic; as I mentioned above, activating these reflective processes is difficult. Not only does the student need to be taught how to appreciate nuance and evaluate evidence, she also has to be encouraged to beware of cognitive laziness and recognize when she needs to override the default automatic system.

III. Types of Reasoning Errors

Now that we have a rough idea of the tripartite model in mind, I can fill out some more of the details. As the reader knows, there are many, many ways a person's thinking can go wrong. In this section, I present a number of examples from the heuristics and biases literature that illustrate the kinds of *systematic* errors that people are prone to make. A closer look at these systematic errors provides empirical support for the tripartite division.

One necessary condition for successfully completing a reasoning task is having the relevant knowledge, what Stanovich calls "mindware." Mindware is the "rules, procedures, and strategies that can be retrieved by the analytic system and used to substitute for the heuristic response" (Stanovich and West 2008: 687). Mindware includes the type of declarative knowledge measured by crystallized g in IQ tests, like vocabulary and general knowledge about the world.7 Mindware also includes methodological knowledge, like probabilistic reasoning, scientific reasoning, and knowledge of formal logical reasoning (Stanovich 2011: 145–46).8

The well-known "Linda problem," created by Tversky and Kahneman, illustrates the importance of mindware. Participants are given the following background information about Linda.

> Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.
> (Tversky and Kahneman 1983: 297)

In one version, the participants were then asked which statement is more probable:

(A.) Linda is a bank teller.
(B.) Linda is a bank teller and is active in the feminist movement. (297)

Tversky and Kahneman tested both statistically naïve undergraduates (who had not had a statistics course) and graduate students who had taken a number of statistics courses. Most statistically naïve participants chose B (85 percent), while only 36 percent of the graduate students chose B (299–300). In post-experimental interviews conducted by Tversky and Kahneman, statistically naïve participants reported that they did not view the conjunction rule (the rule that says satisfying A is more likely than satisfying A and B) as “decisive” (Tversky and Kahneman 1983: 300). It seems that they did not understand how to use the rule to assess probability. This would be an example of what Stanovich calls **mindware failure**; the participants don’t have the relevant understanding of statistical reasoning.

Some of the participants, a majority of the graduate students, do understand the conjunction rule and how to apply it; these students possess the relevant mindware. However, even those with the appropriate mindware can err if they fail to override the automatic system. If one fails to slow down and engage her effortful system, she will respond with the same answer as those who don’t understand the rule. The background information elicits a stereotype (a type of heuristic that guides our automatic judgments) that includes being a feminist, so an automatic response is likely to be B. And, in fact, this is what happened in a version of the Linda problem that asked participants to rank in order of probability eight statements, not just two. In this version, where the choices are not so stark, 85 percent of the graduate students violated the conjunction rule (Tversky and Kahneman 1983: 298). This is an example of what Stanovich calls **failure to override**. In this case, the majority of the graduate students failed to engage the reflective processing system, which, given that they had the relevant mindware, should have led to the correct response.

I don’t want to overstate the case: the efficiency of the algorithmic processes does have an effect on an individual’s ability to reason (Stanovich 2011: 96–97, 249). Once someone has engaged the effortful system, she needs the algorithmic abilities to sustain both the inhibition of the automatic systems and the decoupling of the representation of a thing in the world from the thing itself. To return to the syllogistic reasoning task above, a person must be able to decouple the content in the premises (Living things need water, Roses need water) from what she knows about the world (Roses are living things). To put it slightly differently: successfully answering this question requires the ability to suppress the tendency to endorse the conclusion because one knows independently of the argument that the conclusion is true. Once someone has engaged the reflective system and not just endorsed the automatic judgment, the algorithmic processes come back into play. Continuing to inhibit the automatic processes and sustaining the decoupling are algorithmic processes, not reflective processes, so having a higher IQ will positively affect your ability to reason. But having a high IQ does not entail a tendency to engage the reflective processes: “individuals higher in cognitive ability are more likely to compute the correct response *given* that they have engaged in [effortful] processing, but they are not more likely to actually engage in [effortful] processing” (Stanovich 2011: 148). In other words, having more efficient algorithmic processes (measured by IQ) will help one reason better, but it does not correlate with the ability to recognize when the slow, conscious deliberation of effortful processes is needed or the ability to engage the effortful processes once the need is recognized.9

The absence of even a weak correlation between intelligence and reflective practices shows up in many of the most well documented reasoning errors (Stanovich and West 2008). For example, there seems to be little correlation between IQ and susceptibility to framing effects, anchoring effects or myside bias, three kinds of errors of cognition that have been written about extensively. I'll discuss each in turn, beginning with framing effects. In a 1981 paper, Tversky and Kahneman present a series of experiments illustrating how the way a problem is presented influences our response to it. Here is a famous example from the paper.

> Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill ***600*** people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:
>
> If Program A is adopted, 200 people will be saved. [72 percent of respondents chose this option]
>
> If Program B is adopted, there is 1/3 probability that ***600*** people will be saved, and 2/3 probability that no people will be saved. [28 percent]
>
> Which of the two programs would you favor? (Tversky and Kahneman 1981: 453)

Tversky and Kahneman then asked a different group of people the same question, with the following possible answers.10

> If Program C is adopted 400 people will die. [22 percent]
>
> If Program D is adopted there is 1/3 probability that nobody will die, and 2/3 probability that ***600*** people will die. [78 percent]
>
> Which of the two programs would you favor? (453)

The problem, of course, is that Programs A and C are the same and Programs B and D are the same. People give inconsistent answers, depending on how the issue is framed, in terms of lives saved or in terms of lives lost. The implications of human susceptibility to framing effects are widespread—affecting politics, public policy, and individual life choices.

Another reasoning error with widespread real world implications is the anchor effect (Tversky and Kahneman 1974). The basic idea behind anchoring is that people's decision making procedures can be influenced by factors that have nothing to do with the problem at hand. For example, in one of the earliest demonstrations of the anchor effect, Tversky and Kahneman had subjects spin a wheel, which, unknown to the participants, was rigged to land on the number 65. After they spun the wheel, the participants were asked whether the percentage of African countries that are members of the United Nations was above or below the number spun. A second group of participants went through the same experiment, except that the wheel was rigged to stop on number 10. Despite the fact that the number involved in each of these cases is completely unrelated to the questions about African nations, exposure to the number did affect people's responses. When asked to provide a percentage of the number of African countries in the UN, those who spun a 65 gave higher estimates (45 percent) than those who spun the 10 (25 percent) (Tversky and Kahneman 1974: 1128). One can imagine all kinds of ways people might be exploited and misled by the use of the anchoring effect.

Both framing effects and anchoring effects are classified as tendencies to use a focal bias when processing information. In the case of framing effects, a certain feature of the situation becomes the focus of attention by the way the issue is framed; in the case of anchoring effects the focal bias is produced by the arbitrary number. A third type of focal bias is myside bias, also known as confirmation bias, which is the tendency to evaluate claims in favor of one's own opinion or from one's own perspective.11 For example, Toplak and Stanovich (2003) ran an experiment in which they asked participants to generate arguments on both sides of an issue. (The propositions the participants considered were: people should be allowed to sell their organs, gas prices should be doubled to discourage driving, and university tuition should be paid by the student, not the taxpayer.) In line with other myside bias studies, Toplak and Stanovich found that participants generated more arguments for the side of the issue that they personally supported. They found no correlation between myside bias and IQ. They did, intriguingly, find a correlation with years of schooling; the rate of myside bias decreased with each additional year the participants attended college, even when age and cognitive ability were controlled (856–58). It's not clear what is decreasing the myside bias. In a subsequent study on myside bias Stanovich and West (2007) also failed to find a correlation between myside bias and IQ. But they also found little correlation between thinking dispositions and myside bias.12 One possible explanation of myside bias is that people look for evidence that supports their belief rather than look for evidence that disconfirms their belief. We have a natural tendency to look for evidence that should exist if our belief were true, rather than examining the implications of our belief being false (Klayman and Ha 1987, Wason 1960). It's possible that the decrease in myside bias with more collegiate education is a reflection of the increased tendency to look for disconfirming evidence, but as far as I know, this hasn't been studied. It's also possible that it is the increase of mindware that is responsible for the decrease in myside bias. Most likely, it is a combination of both: students are learning how to look for evidence of alternative hypotheses and they are developing the epistemic dispositions to do so.13

One of the advantages of working with Stanovich's model is that it can help us as teachers locate the points where rationality breaks down. At the most general level, breakdowns occur because of (i) faulty or absent mindware or (ii) because of people's tendency to be cognitive misers. While filling in the mindware gaps is a familiar way to think about pedagogy, counteracting the effects of cognitive laziness is less so. (Though, of course, regular, old laziness is a common complaint.) These general categories can be further analyzed. I examined two frequent mindware gaps, lack of knowledge about probabilistic reasoning (e.g., conjunction errors by naïve undergrads) and ignoring alternative hypotheses when evaluating claims (myside bias). And I looked at several instances of cognitive miserliness: the focal bias effects (framing, anchoring, and myside) and override failure (conjunction errors by sophisticated graduate students).14 Having an awareness of the types of reasoning errors should, I hope, help us in learning how to combat them.

Figure 3. Classification of reasoning errors based on Stanovich's model.

III. Applications in the Classroom

I intend this final section to illustrate ways that my knowledge of Stanovich's work plays out in my own classroom. Though I recognize and respect that there are a variety of ways that his work might inform one's pedagogy, this discussion is based on the kind of teacher I am and the kind of students I teach. For example, I teach a topics-based introductory course and I use a textbook rather than an anthology of primary sources.15 A course relying on primary texts might focus on different kinds of skills and dispositions than the ones I emphasize. And I teach small (around thirty students) classes at a liberal arts school. But I hope that the following discussion will be of general interest to most philosophy teachers. One final disclaimer: I became familiar with Stanovich's work in this area a couple of years ago and incorporating his research into my classroom is still a work in progress.

Before moving into a discussion about my particular course, I'd like to present some of the empirical evidence on the most effective way to incorporate critical thinking dispositions into the classroom. Philip Abrami and colleagues recently published a meta-analysis of the effectiveness of classroom methods on the acquisition of critical thinking skills and dispositions (Abrami et al. 2008). The study looked at four kinds of instruction, general, infusion, immersion, and mixed, based on the work of Robert Ennis (Abrami et al. 2008: 1105–07; Ennis 1989). In a course using *general* instruction, the critical thinking skills and dispositions are themselves learning goals and are taught in a content independent way. In the infusion and immersion courses, content plays a key role. *Infusion* courses include the acquisition of critical thinking skills and dispositions as a learning objective, but the skills and dispositions are not taught independently of the content; *immersion* courses do not make the acquisition of critical thinking skills and dispositions an explicit learning objective. The *mixed* approach teaches critical thinking skills and dispositions as a separate independent track within a course dedicated to a specific content.

Abrami and his colleagues included 117 studies in their analysis (Abrami et al. 2008: 1111). They found a moderate average effect for the instruction of critical thinking skills and dispositions compared to the control group. There is, however, a lot of variability in the data on effect sizes, which suggests that not all instructional methods are equal (1112). A closer look at the data reveals that 32 percent of the variance in effect size can be explained by two factors, the instructional method and the pedagogical training of the instructor (1120). I'm going to ignore the pedagogical training of the instructor factor and focus on the instructional method. Of the four instructional methods discussed above, the mixed method had the greatest effect and the immersion method had the smallest effect. As the authors say, "This is an important finding for the design of courses. Making CT requirements a clear and important part of course design is associated with larger instructional effects. Developing CT skills separately and then applying them to course content explicitly works best" (1121).

My introductory course uses the mixed method described above and is designed with improving critical thinking skills as a primary and explicit learning objective. The first three weeks of my topically structured introductory course are devoted to methodological issues: necessary and sufficient conditions, the nature of philosophical arguments (both deductive and inductive), fallacies, causal and metaphysical possibility, and the structure of thought experiments. The mixed method approach involves teaching the critical thinking skills and dispositions independently of the course content, so I shy away from most philosophical content in the beginning. When introducing the methodological concepts, I use ordinary, familiar topics. For example, when teaching students about necessary and sufficient conditions, I myself use simple, quotidian examples (being colored is necessary for being red, being a German Shepherd is sufficient for being a canine) and ask the students to do the same in their written assignments. When discussing different kinds of fallacies, I ask students to find examples in nonacademic writings and usually get many examples from sports stories or the campus newspaper opinion columns. I bring in philosophical content when we discuss thought experiments, but the emphasis of the assignments inside and outside class is on the structure of the thought experiment. For example, the first thought experiment we discuss in class is the trolley case and the focus of the discussion is how different versions of the trolley scenario bring out conflicting intuitions. The discussion centers around the structure of the thought experiment. I then have the students do an assignment outside class where they are asked to explain what kinds of contradictory intuitions Judith Jarvis Thomson might be trying to elicit in the famous violinist case.16 But I don't let the conversation turn towards broader ethical questions or the rightness or wrongness of abortion; I'm very clear that the point of these exercises is to develop our capacities for philosophical reasoning, not to come to any conclusions about ethics or abortion.

After reading Stanovich, I now think of this period as providing the mindware that the students need in order to succeed in the class. The longer I teach, the bigger this section of the course has become. If they don't acquire the mindware in the beginning, then they won't be able to *do* philosophy as we move through the various topics. In keeping with the mixed method approach, in which critical thinking skills and dispositions are taught separately, I constantly reinforce the importance of the methodological terms. In discussion, I might ask a student to rephrase their point in terms of necessary and sufficient conditions. Or I might ask a student to explain whether her peer's argument is inductive or deductive. And I'm constantly asking the class to explain these basic concepts. For written assessment, I give three exams over the semester and, while the philosophical content changes, each one includes what I call the "logic" section, where I might ask for an explanation of the difference between deductive and inductive arguments or the difference between validity and soundness. I tell students at the beginning of the course that the logic section will be included on all three exams, so they have an incentive to learn the concepts from the beginning.

But providing students with the mindware is not sufficient—you must spend time on teaching them when and how to deploy it. I believe that describing Stanovich's basic model to the students is useful for these tasks, because the model provides a provocative vocabulary and some conceptual distinctions that help to facilitate and encourage metacognition. I explicitly, but informally, discuss dual-processing and Stanovich's addition to the two processes about a third of the way through the class in the section on personal identity and the self. The students read an article on the nature of the self by a psychologist who briefly discusses the dual processing theory. I take that opportunity to talk a little about Stanovich's research. I might not refer explicitly to Stanovich again, but now I can use the concepts of cognitive laziness and reflective processing for the rest of the semester. In addition to giving students some more tools to think about how they think, my brief discussion of pedagogy seems to impress students. They usually appear surprised that my pedagogical approach has empirical foundations (or perhaps any foundations at all).

Until I began this research, I certainly hoped that I was instilling reflective processing habits like looking at different sides of an argument, avoiding dogmatism, thinking carefully before reaching a conclusion, and weighing evidence. But I assumed that these kinds of dispositions would just come from the various readings, activities, and discussions we pursued, or maybe from modeling good epistemic behavior. But reading Stanovich impressed upon me the cognitive difficulty of engaging the reflective processes and overriding the automatic system. And the meta-analysis by Abrami et al. reinforces the idea that my passive approach was likely ineffective: "Developing CT skills separately and then applying them to course content explicitly works best; immersing students in thought-provoking subject matter instruction without explicit use of CT principles was least effective" (Abrami et al. 2008: 1121).

So I am trying to be much more intentional in how I teach critical thinking skills and dispositions. Here is the kind of activity I do with my students to illustrate the myside bias, which also demonstrates how I work in the mindware with philosophical content. Say we are discussing Singer's "Famine, Affluence and Morality." I'll ask the students to bring to class a reconstruction of his main argument in numbered premise form. They'll meet in small groups and compare their versions of the argument and then I'll go over the argument with the whole class to make sure everyone understands what the main argument is and how we got there. Once everyone seems to grasp the argument, the students are told to work for a defined amount of time (five to ten minutes) individually or in pairs to come up with a list of objections and a list of supporting claims for Singer's conclusion. I then ask people to share the number of objections or supporting claims they came up with and ask them to compare these to their own views. If they disagreed with Singer's conclusion, did they find it easier to come up with objections? I use this as an opportunity to explain what myside bias is and how pervasive it can be, before moving back to a discussion of Singer's argument. I'll keep bringing up myside bias the rest of the semester. In particular, we discuss it again when I hand out the final paper topics, which requires the students to defend or prosecute (their choice) a man with a brain tumor who is on trial for murder. The paper assignment requires that they consider well-developed objections from the other side. In discussing this part of the assignment, I again initiate a discussion about myside bias and how difficult it can be to develop arguments that challenge one's beliefs.

The framing effect is another kind of bias that can be easily worked into philosophical discussions. When I bring up the trolley problem again, this time in a philosophical rather than a methodological context, I present both the original version and the pushing-a-man-off-the-bridge version and we discuss how an arguably morally irrelevant difference changes our intuitions. I also bring up framing effects during the unit on personal identity. Another clear example is Joshua Knobe's (2003) chairman case, in which the side effects of an agent's action are deemed more intentional if the side effects are bad than if the side effects are good. In both scenarios, the agent's primary action is the same and performed for the same reason; the only difference is whether one of the side effects of the action is to harm or help the environment. I haven't used Knobe's case yet in my introductory class (though I have used it in my epistemology course) but I plan to incorporate his case and the Tversky and Kahneman Asian flu case into a short lesson plan. I want to reinforce the notion that these kinds of effects occur outside the study of philosophical thought experiments.

All teachers struggle with getting students to recognize that the skills they learn in class are transferable to other domains. Incorporating this understanding of how we reason into the classroom might help with this problem. More specifically, Stanovich (2009) mentions the examples of John Allan Poulos, the famous mathematician, and David Denby, the movie critic for *The New Yorker*, as illustrations of smart, accomplished people making incredibly bad financial decisions. These particular examples could be useful in the classroom, since Poulos and Denby have written about their follies (Poulos 2003, Denby 2004).17I haven't done this yet, but next semester I plan to use the final week of class to focus on the theme of smart people doing stupid things,which will give me one more chance to emphasize the importance of developing our reflective processes.

Framing the course in this way—where a main learning objective is to teach students to reflect on their deliberative processes—has had an impact, at least in self-reports.18 Many students believe that they are more aware of their own biases. I frequently get feedback that the skills and dispositions we work on help in other courses they take, both in discussion and in written work. I am lucky to have a colleague who has spent the past few years developing a department assessment plan for how to measure thinking dispositions, and I hope to work with him individually in the future to get beyond these self-reports.19 I have found the empirical work I presented helpful both in improving my general approach to teaching reasoning skills and in developing specific kinds of content. My main purpose in writing this paper, however, was to convey what I believe is valuable information about human cognition; I trust that you can figure out how best to use it in your own classrooms.20

Notes

1. In the literature, these are called Type I and Type II or System I and System II.
For the reader not familiar with this literature, I thought descriptive labels might be more useful.
2. This is not always a bad thing. Relying on heuristics is often quite useful. See Gigerenzer, Todd, and ABC Group 2000.
3. IQ scores have been rising ever since they were first tested. It turns out that this is due to rising Gf scores. See Flynn 2007 for discussion.
4. I'm not arguing that "intelligence" should be generally used in this narrow way, but merely defining the term operationally.
5. Stanovich and West (this chapter was co-written) provide a table of the tasks and effects that are and are not correlated with IQ. I focus on those that are not correlated, since I think they should be a special focus of teachers. Some of the effects that are correlated with intelligence are: outcome bias, hindsight bias, covariation detection, and the overconfidence effect. For a complete list, see Stanovich 2011: 140.
6. See Stanovich 2011: 139–45 and Stanovich and West 2008 for discussion of what tasks do and do not correlate with intelligence.
7. Here is a depressing example of a lack of declarative knowledge. In the textbook I use, one problem asks for the validity and soundness of an argument with a premise stating Mexico City is south of Bogota. Very few of my students know where Bogota is and thus most are unable to evaluate the claim of soundness.
8. Serious philosophical questions can be raised about how sharp the distinction is between the various types of processes. The methodological knowledge that Stanovich is here categorizing as

mindware seems likely to involve the dispositions that Stanovich thinks characterizes the reflective processes. For more critical discussion, see Woods 2012.

9. It is probably worthwhile to keep in mind that the great majority of the participants in these experiments, like in most psychology experiments, are college students, a population with an above average IQ. So it is possible, though it seems to me unlikely, that with a sample that better reflected the normal distribution of IQ there might be a correlation between IQ and these reflective processes. Since most readers of this article are likely to be teaching college students, I am going to ignore this potential complication.

10. Let me emphasize that this is a between subjects design, which means that participants saw only one of the versions. Unless noted, all experiments discussed are between subjects design. A within subjects design (where the subject receives multiple versions of a problem) is likely to alert the subject to a conflict in her beliefs. An experimental design in which the subject is alerted to an inconsistency in her belief set would be measuring maximal performance, rather than typical performance.

11. It is important to distinguish myside bias from belief bias. Belief bias was illustrated in the syllogism example—people are likely to assent to the validity of the syllogism because they have prior knowledge that roses need water. Myside bias involves people's opinions, like the belief that the gas tax should be raised. While no correlation has been found between IQ and myside bias, there is a correlation between IQ and belief bias. See MacPherson and Stanovich 2007 for details.

12. There is evidence that suggests that myside bias is domain specific. A person's myside bias in a given domain does not predict myside bias in other domains. So, for example, a person might show myside bias when confronted with questions about tuition subsidies, but not when asked about organ selling (Toplak and Stanovich 2003, Stanovich and West 2007).

One might worry that the findings of Stanovich and West (2007) challenge my thesis that we should focus more of our attention on reflective processing. Here are two ways to temper that worry. First, our psychometrics for reflective processing is less developed that our psychometrics for intelligence. (See Stanovich 2011: 191–246 for discussion of what a Rationality Quotient [RQ] test might look like.) Stanovich and West (2007) used two different measurements for reflective processing, but there might be other kinds of testing that do correlate with decreased myside bias. Second, as noted in the main text, myside bias does decrease with additional years of college. Since students are gaining both mindware and new thinking dispositions, I think it likely that both are contributing to the decrease.

13. Thanks to an anonymous reviewer for pushing me to develop this point.

14. Here is another example of where the categories seem loose. Myside bias can be categorized as both a mindware failure and a result of cognitive miserliness. I don't, however, see this as a problem. One can think of the search for alternative possibilities as a method that needs to be taught (mindware) and also recognize that the automatic system must be overridden in order to deploy this method. I think that the conjunction fallacy should be similarly understood.

15. I recognize this is appalling to many philosophy teachers, but that is the topic of another paper.

16. If you are wondering, the use of such a controversial topic is intentional.

17. Edmund L. Andrews, an economics reporter for the *New York Times*, has also written an article on the topic, "My Personal Credit Crisis," *New York Times Magazine*, May 14, 2009.

18. At the end of the semester, I ask the students to write a brief reflective essay on the course that includes discussing the most important thing they learned. About a third of the class usually responds with methodological issues and this percentage has increased as I've spent more time talking with them about my goals for the course.
19. See Wright and Lauer 2012 for a detailed discussion of their assessment efforts.
20. Thanks to Jean Lavigne and two anonymous reviewers for their thoughtful comments on an earlier draft of this article.

Bibliography

Abrami, Philip C., Robert M. Bernard, Evgueni Borokhovski, Anne Wade, Michael A. Surkes, Rana Tamim, and Dai Zhang. 2008. "Instructional Interventions Affecting Critical Thinking Skills and Dispositions: A Stage 1 Meta-Analysis," *Review of Educational Research* 78(4): 1102–34. http://dx.doi.org/10.3102/0034654308326084

Denby, David. 2004. *American Sucker*. Boston: Little, Brown and Company.
Ennis, Robert H. 1989. "Critical Thinking and Subject Specificity: Clarification and Needed Research," *Educational Researcher* 18(3): 4–10. Available at http://www .jstor.org/stable/1174885.

Facione, Peter A. 2000. "The Disposition Toward Critical Thinking: Its Character, Measurement, and Relationship to Critical Thinking Skill," *Informal Logic* 20(1): 61–84.

Flynn, James R. 2007. *What is Intelligence?: Beyond the Flynn Effect.* Cambridge: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511605253

Gigerenzer, Gerd, Peter Todd, and ABC Group. 2000. *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Hassin, Ran R., James S. Uleman and John A. Bargh, eds. 2006. *The New Unconscious*. New York: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780195307696.001.0001

Kahneman, Daniel. 2011. *Thinking Fast and Slow.* New York: Farrar, Straus, and Giroux.

Klaczynski, Paul A., and Kristen L. Lavallee. 2005. "Domain-Specific Identity, Epistemic Regulation, and Intellectual Ability as Predictors of Belief-Based Reasoning: A Dual124 Process Perspective." *Journal of Experimental Child Psychology* 92 (September): 1−24. http://dx.doi.org/10.1016/j.jecp.2005.05.001

Klayman, Joshua, and Young-Won Ha. 1987. "Confirmation, Disconfirmation, and Information in Hypothesis Testing," *Psychological Review* 94(2): 211–28. http://dx.doi.org/10.1037/0033-295X.94.2.211

Knobe, Joshua. 2003. "Intentional Action and Side Effects in Ordinary Language." *Analysis* 63: 190–93. http://dx.doi.org/10.1093/analys/63.3.190

MacPherson, Robyn, and Keith E. Stanovich. 2007. "Cognitive Ability, Thinking Dispositions, and Instructional Set as Predictors of Critical Thinking," *Learning and Individual Differences* 17: 115–27. http://dx.doi.org/10.1016/j.lindif.2007.05.003

Poulos, John Allen. 2003. *A Mathematician Plays the Stock Market*. New York: Basic Books.

Singer, Peter. 1972 "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1(3): 229–43. http://www.jstor.org/stable/2265052.

Stanovich, Keith E. 2009. *What Intelligence Tests Miss: The Psychology of Rational Thought.* New Haven, Conn.: Yale University Press.

. 2011. *Rationality and the Reflective Mind.* New York: Oxford University Press.

Stanovich, Keith E., and Richard F. West. 2007. "Natural Myside Bias is Independent of Cognitive Ability." *Thinking and Reasoning* 13(3): 225–47. http://dx.doi.org/10.1080/13546780600780796

. 2008. "On the Relative Independence of Thinking Biases and Cognitive Ability." *Journal of Personality and Social Psychology* 94(4): 672–95. http://dx.doi.org/10.1037/0022-3514.94.4.672

Thomson, Judith Jarvis. 1971. "A Defense of Abortion," *Philosophy and Public Affairs* 1(1): 47–66. Available at http://www.jstor.org/stable/2265091.

Toplak, Maggie E., and Keith E. Stanovich. 2003. "Associations Between Myside Bias on an Informal Reasoning Task and Amount of Post-Secondary Education." *Applied Cognitive Psychology* 17: 851–60. http://dx.doi.org/10.1002/acp.915

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases," *Science* (New Series) 185 (September 27): 1124–31. http://dx.doi.org/10.1126/science.185.4157.1124

. 1981. "The Framing of Decisions and the Psychology of Choice," *Science* (New Series) 211 (January 30): 453–58. http://dx.doi.org/10.1126/science.7455683

. 1983. "Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," *Psychological Review* 90(4): 293–315. http://dx.doi.org/10.1037/0033-295X.90.4.293

Wason, Peter. 1960. "On the Failure to Eliminate Hypotheses in a Conceptual Task," *Quarterly Journal of Experimental Psychology* 12: 129–40. http://dx.doi.org/10.1080/17470216008416717

Woods, Cathal. 2012. "Book Review: Rationality and the Reflective Mind," *Informal Logic* 32(3): 373–80.

Wright, Charles W., and Abraham Lauer. 2012. "Measuring the Sublime: Assessing Student Learning in Philosophy," *Teaching Philosophy* 35(4): 383–409.
http://dx.doi.org/10.5840/teachphil201235443