

# Critical thinking and debiasing<sup>1</sup>

Tim Kenyon, Department of Philosophy, University of Waterloo

Guillaume Beaulac, Rotman Institute of Philosophy, University of Western Ontario

## I. Introduction

Teaching critical thinking is an undertaking that permits emphasis on many different combinations of elements, the most traditional of which are formal logic, informal logic, argumentation, fallacy theory, and rhetoric. Increasingly, however, critical thinking courses and texts also include an explicit emphasis on the psychology of cognitive and social biases (see, for example, Kenyon 2008; Ruggiero 2004; Groarke & Tindale 2004; and Gilovich 1991). The evident aim is to enable students to identify biases in reasoning, and to minimize biases in their own thinking, by informing them about character and contextual features of those biases.

This is a sensible aim in itself, consonant with the general rationale for teaching critical reasoning courses in the first place. But it presupposes that simply teaching students about biases is an effective way of enabling students to reduce the distortions of biases in their own thinking. We identify this assumption as the *intuitive approach to teaching debiasing*.

---

<sup>1</sup> This is an unpublished version; we ask that it not be cited without our permission. Correspondence regarding permission should be directed to: Tim Kenyon, Department of Philosophy, University of Waterloo (tkenyon@uwaterloo.ca). We would like to acknowledge the helpful input of Frédéric-

(IATD) Teaching facts about biases, including a taxonomy of biases and their various propensities to distort reasoning, is reasonably efficacious as a means of imparting debiasing skills to students by way of critical reasoning courses.

When one considers the empirical evidence bearing on that question, however, the most plausible simple answer is that IATD is false. At least, the practice of simply teaching students facts about biases is not as effective as one might hope. The literature on the psychology of debiasing indicates, on balance, that teaching people about biases doesn't reliably debias them. Indeed, that literature suggests that (for at least a wide class of biases) practically *any* debiasing strategy intended to be learned and subsequently self-deployed by individuals acting alone, at the point of making a judgment, is unlikely to succeed in significantly minimizing biases.

In the following remarks, we briefly outline why this is so before moving on to consider the ramifications for critical thinking education. Vast resources are currently devoted to teaching critical reasoning worldwide. Does the implausibility of IATD mean that these resources are misused, to the extent that they are predicated on IATD? Should philosophers, psychologists, and other critical reasoning educators just stop including a focus on biases in critical thinking education?

We don't think so. Rather, we take the lesson to be that whole societies and polities have a major interest in promoting efficacious debiasing education – extending to population-level demographic scales and intergenerational time scales.

The difficulty of teaching debiasing skills that can be deployed in a strictly atomistic or individualistic way counts in favour of teaching and investing also in more collective debiasing strategies and infrastructure that would serve the latter sorts of interests. This approach will encompass teaching not just individual skills and knowledge, but skills that enable the construction of reasoning infrastructure, and effective participation in social and organizational reasoning processes and decision procedures.

What would these processes, strategies and infrastructure look like? A key first step here is to reflect on the breadth of what can count as debiasing from a critical thinking perspective. Our aim in this reflection is to help motivate and set the stage for creative and empirically guided work on how to teach debiasing in ways that might be efficacious, serving both private and public interests in minimizing distorted or unreliable reasoning. By focusing on choices, behavior, and agent-world interactions, we suggest a broader range of outcomes for critical thinking than that informing IATD, and therefore a broader range of options for critical thinking education as well.

## II. Pessimism about teaching debiasing abilities

First, a caveat. The empirical literature on biases and on debiasing is massive and varied; even to summarize it comprehensively would be impossible for a single paper. We will use just a few representative results to illustrate the grounds for thinking that teaching students about biases, and warning them to be on the lookout

for biases, is unlikely to significantly reduce the generation of distorted judgments “in the wild,” or to increase the likelihood that biased judgments will be recognized and remedied by the agent herself. Because we hope to spend some time on the *implications* of this fact for critical thinking education, we are compelled to move through the empirical evidence rather briskly. So our subsequent reflections will have to remain conditional, not just on the probity of the defeasible empirical literature, but on the accuracy of our depiction of that literature<sup>2</sup>.

The most general problem to emphasize is illustrated in Baruch Fischhoff's (1982) seminal work on mitigating the *hindsight bias*. Hindsight bias is the tendency to regard actual outcomes as inevitable outcomes, in retrospect, and to overestimate the extent to which one had antecedently expected the actual outcome. Fischhoff points out that a number of common-sense approaches to debiasing subjects for hindsight effects simply do not work very well under a wide range of experimental conditions (1982, pp. 427-431). These approaches include:

- Explicitly explaining the bias to subjects, and asking them to avoid it in their own reasoning;
- Inducing subjects to value the accuracy of their performance;
- Encouraging subjects to think first in terms of diagnosing other subjects' biased reasoning, before turning to the question of their own reasoning.

---

<sup>2</sup> A somewhat more detailed examination of this evidence is provided in Kenyon 2014.

These strategies (and others that Fischhoff describes) are motivated by some quite natural assumptions about the nature of learning, of cognition, and of error – the most basic one being that forewarned is forearmed. But their ineffectiveness at mitigating hindsight bias, and many other biases, has been quite strongly confirmed by subsequent psychological work (Wilson et al. 2002).

Fischhoff's studies and subsequent ones are, perforce, largely experimental designs that isolate specific instances of biased reasoning, rather than longitudinal analyses of learning outcomes in educational contexts. The latter type of study, though, tends to be hard to perform rigorously, and hard to interpret (Willingham 2007, p. 12). Experimental designs constitute the best evidence we now possess about the propensity for normally teachable information and skills to reduce biased reasoning in students, outside the classroom and in later life. How much confidence to place in the applicability of these results is a good question; but this is the evidence we have. On balance, it weighs against the thought that simply teaching and warning people about biases will successfully mitigate biased reasoning. IATD is not well supported by evidence.

Perhaps the most significant factor explaining why teaching people about biases doesn't itself particularly reduce their biases is known as *bias blind spot* (Pronin and Kugler 2007; Pronin, Lin & Ross 2002). Put simply: knowing that people in general are subject to a particular bias is consistent with one's believing that one is not subject to it. Indeed, more importantly, even knowing that one is *generally* susceptible to a bias is consistent with one's believing, on the specific occasion one considers the matter, that one is not displaying a bias. Bias blind spot thereby

insulates one's judgments, in the event, from the application of whatever debiasing strategies might actually be effective.

A related explanation for the relative ineffectiveness of teaching information about biases is that we can easily think that we have debiased when we have not. A theoretical knowledge of the need to adjust for bias does not reduce this problem, since the problem is precisely that one falsely believes oneself to have addressed that need. Indeed, merely thinking about debiasing can enable the problem! By thinking over the details of the case at hand, and considering the prospect of being biased, one may simply give one's biases more raw material to operate on (Thompson 1995). Here, then, another appealing thought about clear thinking meets unwelcome data: the idea that one can debias by firmly *thinking it over*, that debiasing can be a matter of having a stern word with oneself about not being biased, is mistaken (Frantz & Janoff-Bulman 2000). In fact, attempting to self-debias in this way can even make one's biases worse (Hirt and Markman 1995; Sanna, Stocker & Schwartz 2002). Telling ourselves that we have debiased, we can come to hold our attitudes and views more strongly – convinced that they have been vetted for distortion. As Frantz (2006, p. 165). observed, merely to ask a question like “Am I being fair?” is to provide an additional opportunity for a bias to operate, accompanied by a greater conviction that the answer is “Yes.”

This conveys a sense of the kinds of evidence speaking against the idea that we can teach people to be significantly less biased reasoners simply by teaching and warning them about biases. But this isn't to say that *no* debiasing strategies have been shown to work in this literature. A range of strategies work to varying degrees,

depending on the bias, with the single most effective (and most *generally* effective) strategy being for the subject to explicitly consider and entertain a range of alternative perspectives or counterfactual outcomes, and what would have had to happen in order for those outcomes to occur (Pronin, Puccio and Ross 2002; Wilson et al 2002; Anderson & Sechler 1986; Fischhoff 1982). So we do have at least one mitigation strategy with a significant prospect of success, *taken as an experimental treatment*.

The problem is that the strategy is extremely difficult to implement as a self-deployed skill. Existing biases and attentional limits can easily make themselves felt as an unwillingness or inability to generate plausible alternative scenarios (O'Brien 2009, pp. 329-330); and even a willingness to do so is no guarantee that the generation and consideration of alternatives will be sufficiently disciplined or constrained to actually lead to a less distorted judgment (Tetlock 2005, p. 199). Absent the facilitation or guidance by assistants that tends to characterize experiment contexts in which "consider the opposite" is an effective strategy, there is little reason to expect it to be employed with regularity by individual agents, nor to work well when it is employed.

Roughly and readily, then, there is a seeming dilemma for those who wish to teach debiasing as part of critical thinking. The things that are most easily teachable and open to long-term retention by learners – what biases are and how they work; and that biases are to be avoided – are not in themselves very effective at debiasing people's judgments; while the things that are rather effective at debiasing judgments – counterfactual or opposite-scenario consideration – aren't very

teachable as individual skills to be recalled and applied when needed, nor to be implemented easily even when attempted.

Again, we do not take this as grounds to doubt that there is still a rationale for focusing on biases in critical thinking education. Rather, we believe the problem of teaching effective debiasing strategies is really just an invitation to a more fine-grained taxonomy of debiasing outcomes. If teaching debiasing looks too hard in light of these data, it is because IATD focuses only on doing it at the least plausible levels: by creating debiased people, or even just debiased *thinking*, in the event. We propose to mitigate this problem by distinguishing further levels or domains of debiasing. This in turn should open up a wider range of overlapping skills and habits that can more plausibly be taught and implemented, aimed at addressing biases at those different levels.

### III. The scope of debiasing

While knowing about a bias is no prophylactic in itself, it may serve as one of many steps along a path to debiasing (Stanovich & West 2008; Wilson & Brekke 1994). For example, Wilson & Brekke's model lists the awareness of an unwanted process as the first step in debiasing. Of course, one must also be motivated to correct the bias, know the *direction* and the *magnitude* of this bias, and be sufficiently in control, with sufficient mental resources, to be able to adjust the response (Wilson & Brekke 1994, p. 119). Now, we have reviewed grounds to believe that no individually portable suite of skills seems very apt to put these cognitive and affective resources



to work at the right times. But what if we did not limit ourselves to the goal of preventing biased judgments, nor even that of unskewing judgments after they are made?

We take the problem thus far to be an artefact of the level at which we have been considering both biases and debiasing strategies. We submit that the core issues of interest from a critical thinking perspective are broader – including not simply what one thinks, but how one acts.<sup>3</sup> This opens up the scope of what will count as a debiasing strategy in the relevant sense. It holds out the promise that a more variegated conception of bias-reduction will offer a range of strategies that limit bias at different levels and in different ways.

In effect we propose swapping a seemingly intractable teaching problem for what is plausibly a more tractable one. The intractable problem is that of teaching the skills of avoiding bias blind spot and implementing debiasing strategies so as not to simply recapitulate the bias. The tractable problem is that of teaching the habits and skills of adopting *general* action and decision-making principles that both nudge agents away from biased reasoning and filter its effects out of their actions.

When we talk of a *nudge*, we mean the term in the sense advanced by Thaler & Sunstein (2009). A nudge is a strategy or a bit of infrastructure put in place in order to minimize or to eliminate a set of cognitive biases by using aspects of the environment. Common examples of nudges include changing the way information is presented to participants in a process, and changing the default options in some task. A striking example given at the outset of Thaler & Sunstein's discussion is the

---

<sup>3</sup> C.f. Beaulac & Robert (2011) on critical thinking *attitudes*.

way food is displayed in a buffet: depending on where certain food items are placed, their popularity as a choice may rise by 50%. The idea, then, is to pre-emptively construct situations in order to minimize biases. From an individual agent's perspective, this presents two dimensions of control: whether to exploit existing nudges in the environment, and whether to construct nudges of one's own – either individually or collaboratively. Both the ability and the need to do these things generally are potential learning outcomes for a critical thinking course.

While a more fine-grained analysis is surely possible, we will for the sake of brevity limit the current discussion to four broad levels at which debiasing can be implemented, once it is taken to span the distinction between thought and action. We provide both a general description and an example for each level.

Level 1 debiasing: Owing to critical thinking education and training, an agent has no disposition to produce a particular sort of biased judgment; that is, the bias does not arise. This sort of debiasing process is implemented during education and applies to individual agents' judgments.

E.g., A hiring committee member does not notice or attend to racial differences, and shows no bias in reasoning about the quality of candidates from visible minority groups in hiring context.

Level 2 debiasing: A biased judgment occurs or is incipient, but critical thinking education and training facilitate the agent's deployment of

cognitive or behavioral strategies that lead to a revision of the judgment in context. Debiasing of this kind is implemented within the context of judgment-fixation, is initiated and mediated by agents' psychological processes, and applies to individual agents' judgments.

E.g., A hiring committee member's first reaction is to assign an unwarrantedly low rating to a dossier from a candidate with a name connoting ethnic minority status. On second thought, though, she wonders whether she is being biased by the character of the name, and reflects on the positive features of the file. Eventually she comes to think of the candidate in more accurate terms.

Level 3 debiasing: A biased judgment occurs or is incipient, but critical thinking education and training (individual or collective) leads (or has led) to the creation of situational "nudges" that debias the agent's judgment in context. This sort of debiasing process is implemented within the context of judgment-fixation, is initiated or mediated by environmental cues or infrastructure, and applies to individual agents' judgments.

E.g., A hiring committee is given a preliminary presentation about the prospects for biased reasoning in hiring contexts. Notes and other guidelines from this presentation are kept in the meeting room, in a red folder on the table around which committee members sit. Later, a hiring committee member encounters a dossier from a candidate with a name connoting ethnic minority status. The visual salience of the red folder

reminds her to attend to the significance of the candidate's name. She would otherwise have assigned an unwarrantedly low rating to the file, but owing to the earlier presentation she makes a point of reflecting on the candidate's positive features, considers how those features would appear if part of a privileged candidate's application, and ranks the file more accurately.

Level 4 debiasing: A biased judgment occurs, and is not significantly remedied, but situational constraints nevertheless debias the action. This type of debiasing process is implemented over time, both in advance of and during the context of judgment-fixation. It is initiated or mediated by environmental cues or infrastructure, and applies to group judgments, or to actions and outcomes.

E.g., A hiring committee member has an uncorrected bias of judgment against women in the profession; but anonymized applications hide candidates' gender information, and the committee member ultimately (unknowingly) votes to hire a superior woman candidate.

E.g., A hiring committee member displays uncorrected biased reasoning in judging that a superior candidate should not be hired because of her sexual orientation; but declines voice this view in light of the negative responses it would draw from colleagues, and ultimately votes in favor of the candidate.

E.g., A hiring committee member displays uncorrected biased reasoning in judging that a superior candidate has an inferior track record, but the majority vote of the hiring committee favours the candidate, and she is offered the job anyhow.

The levels represent a way of carving up the gradient from the most individualist and internalist character-driven approaches, to the most outcomes-oriented and externally-mediated approaches. We can characterize Levels 1 and 2 as the more individualistic levels; they essentially treat the particular agent as both the source and the focus of debiasing outcomes. Levels 3 and 4 appeal to external, situational factors to a greater extent.

Level 3 debiasing retains a crucial individualistic component, since the “nudges” or external aids to reasoning that it postulates are devoted to mitigating biases in the individual agent. With Level 4 debiasing, this individual aspect is minimized, in some cases to the point of being eliminated altogether. The external factors it invokes are essentially oriented towards debiasing decisions, actions, and outcomes – including group outcomes – without specific reference to the dispositional properties of any particular agent. Level 4 debiasing will of course still have individualistic overtones, since an agent may learn in retrospect what and how to think about an issue by seeing a debiased outcome. Indeed, this may well be a valued feature of such debiasing efforts and infrastructure over the longer term. But it is not a defining feature of Level 4 debiasing success.

Our conjecture is that, when it comes to biases, IATD and similar naïve approaches to critical thinking education treat Level 1 outcomes as the ideal, and strive at least to bring about Level 2 outcomes. We think this is practically impossible; if such education is ever effective, it is more likely because elements of the education itself are acting as persistent nudges to create occasional Level 3 outcomes, while the value of Level 4 outcomes is learned by trial and error, if at all, and is implemented relatively haphazardly. The impetus to treat Levels 1 and 2 as the real aim of critical thinking education depends, we think, not on evidence that this is a practical possibility, but substantially on a deep-seated intuition that critical thinking is properly implemented only in the minds and choices of specific agents.

The three distinct examples for Level 4 debiasing reflect both the flexibility of the individuation of actions, and the range of points at which debiasing action can take place. The first example proposes an intervention affecting the agent's judgment of candidates; with the second, the intervention debiases the agent's act of voting; while the third describes a mitigation of nothing more specific than the committee's collective hiring actions. The anonymized hiring protocol, the perception of social disapproval of prejudice, and the committee voting structure each count as an element of contextual engineering that effectively debiases the Level 4 scenario, even though all Level 4 cases *by definition* count as failures of debiasing by the purely individualistic cognitive standards we originally considered.

Clearly, then, this more fine-grained analysis reveals more opportunities to debias by clarifying the number of stages open to intervention in thinking, preparing, deciding and acting. Variations on the theme are not hard to find,

moreover, including some that span the levels we have sketched. For example, Uhlmann and Cohen (2005) found that, if the notion of merit were left undefined for a hiring process, it would tend to become the vehicle of gender-biased decision making. That is, merit would be operationalized distinctly from case to case, with the overall effect of promoting hiring along gender lines – and particularly the hiring of men over women<sup>4</sup>. But eliciting a commitment to some hallmarks of merit from the evaluators *prior* to revealing information about the people being evaluated reduced this biased “moving goalposts” approach in their judgments (2005, p. 478). The example provides further empirical support for the idea that education about the prior construction and acceptance of such policies and organizational structures should fall within the core mandate of education for reaching more appropriately reflective and reliable outcomes in reasoning.

Arguably this counts as a remedy that straddles the border between Levels 3 and 4, since the incipient bias is corrected in judgment, not merely in action or outcome; yet in practice the mechanisms achieving this outcome will be thoroughly environmental and causally remote. That is, somebody has to decide (presumably well in advance, in the case of policy-making) to set out clear rubrics for merit, and to ensure a hiring process structured so that evaluators review the hallmarks of merit before they review the details of applicants. So not every case of debiasing falls entirely within one such level; but we do think that this particular way of

---

<sup>4</sup> There was also weak evidence that female evaluators would similarly construct merit in a gender-biased way to devalue male applicants, if the job were sufficiently stereotypically associated with women’s gender roles – e.g., that of a Women’s Studies professor (2005, p. 478).

carving up of levels helps illuminate relevant features of even those cases spanning levels.

#### IV. Teaching debiasing as teaching the acceptance of influences on cognition and of constraints on action

There is, then, a very broad recipe for achieving better odds of teaching successful debiasing strategies: first broaden the conception of what counts as debiasing, and then be open to exploiting the full spectrum of opportunities to mitigate bias, from antecedent reasoning dispositions to the broadest conception of an action in context. We close with some schematic remarks about putting debiasing, so construed, into a typical critical thinking curriculum.

Critical thinking pedagogy, we believe, becomes more effective than the debiasing strategies described earlier when it motivates us to subject ourselves to nudges, infrastructure and institutions in advance of the circumstances of bias that will make those things effective debiasing aids. That is, we suggest that knowledge of biases has the best chance of effectiveness when it leads one generally to accept and construct nudges or contextual engineering of one's own. In that case it supports the adoption of general debiasing strategies that might simply be encoded in the lived environment, rather than holding out the hope that one can learn to debias in a series of contextual one-offs, as the need arises.

An education in debiasing therefore ought to include information and skill-development regarding the administration of decision-making contexts and actions,



in a manner consonant with Level 3 and Level 4 debiasing. A rough analogy might hold with remedial education for a person addicted to gambling: rather than trying to teach the ability to avoid gambling while at the casino, we would counsel teaching the ability to avoid the casino. Choices ranging from how to strike committees, how to solicit information, which buttons to push on the television remote control, and whether to ask about someone's personal details can all powerfully influence the opportunities for biases to be reflected in our actions. It follows that the knowledge (both knowledge-that and knowledge-how) associated with those activities are reasonable components of an education in critical thinking. This knowledge will include skills of creating and maintaining physical, institutional and social infrastructure that facilitates more truth-conducive reasoning. But often this infrastructure already exists when students and former students encounter contexts of judgment and action; in those cases, the relevant skill will be that of deferring to such truth-conducive mechanisms. How to teach this knowledge and these action principles is a good question. Its feasibility, though, seems far more promising than the mere hope of IATD, that some combination of knowledge of biases and mental continence will be both effective and learnable.

It is worth considering a potential objection to the Level 4 style of debiasing education proposed here, proceeding from an amalgam of epistemological and pedagogical scruples. The worry is this: whatever the didactic barriers to focusing on Levels 1, 2 and 3 as debiasing strategies, addressing one's teaching to these levels at least promotes the right connection between methods and outcomes. By teaching students to recognize bias-inducing situations and to mobilize appropriate

debiasing strategies in context as individuals, one would be teaching students to make cogent *inferences* regarding the need for unbiased or less-biased reasoning. Level 4's blunt focus on debiased *outcomes* does not require anyone in the context to appreciate the problem, nor why it is a problem, nor how the debiasing mechanisms will address the problem. For all that a Level 4 approach tells us, successful debiasing processes can be entirely arational from the perspective of the agents in the situation. But that misses something valuable about the rational connections between reasons and outcomes, something that people should be taught to entertain, not to elide. How can this be a critical thinking strategy when it encompasses solutions that do not involve thinking about the problem at all?

The worry is based on an overly narrow conception both of the scope of the problem and of the problem-solving context. Here it may be useful to return to the analogy with an addict who avoid pathological activities by avoiding situations that lead to those activities. The gambling addict need not avoid gambling situations solely by reflecting on the evils of gambling, nor need she choose to do some other activity on the basis of such reflections *at the time of the engaging in the alternative activity*. She might come to engage in a non-gambling activity out of sheer habit; but if she originally cultivated that habit as a means of avoiding gambling, then any particular case of avoidance by way of that activity reflects her considered judgment and her autonomy.

The proposal at hand articulates a similar (minimally sufficient) connection between education about biases and students' subsequent participation, possibly just from habit, in cognitive and social routines and practices that promote reliable

reasoning. A rational and appropriately agent-endorsed connection between outcomes and methods is established when students are educated about the need to form such habits, or to defer to truth-conducive judgment and action mechanisms. Subsequently acting on those habits need not itself be an exercise in reasoning or inference at the point of action in order to be an exercise of the agent's commitment to critically informed reasoning. Indeed, consciously reviewing one's reasons at the point of decision-making might even disrupt the debiasing process. To deny that this represents the exercise of critical reasoning is to deny, by parity of reasoning, that the lifelong alcoholic who cultivates a preference for badminton as an alternative to hanging out in a pub is not demonstrating a willful continence regarding alcohol when she remains sober for years on end by spending time at the gym.

Examples that we take to be of a similar shape are already implemented in some institutions, with the case of anonymized musical auditions being particularly telling. Women have long been underrepresented in orchestras around the world, comprising fewer than 10% of musicians in major American orchestras prior to the 1970s and little more than 20% in the 1980s, a much lower proportion than their availability in the hiring "pipeline" (Goldin & Rouse 2000). Of the various practices introduced by orchestras to reduce biases that might account for this imbalance, the most common means is to ensure the anonymity of candidates during auditions, by placing the musician behind a screen where he or she plays for 5 to 10 minutes (Goldin & Rouse 2000, p. 722). The screen is not used uniformly across orchestras; only three of the 11 orchestras discussed in by Goldin & Rouse employ it all the way through the process (2000, p. 723). The effects of the screen, however, are

remarkable: when the screen was used throughout the process, the probability that a woman would be offered the job was 60% higher than without it.

The use of the screen is clearly a Level 4 debiasing strategy in our taxonomy: its success in debiasing the outcome of the decision process does not require reduction of the dispositional or occurrent biases of the individual deciders at the point of evaluating candidates. Yet the prior decision to implement a general policy of anonymized auditioning is plausibly driven by just the sort of empirical details about biases, and commitments to erring on the side of caution, that a sound critical thinking education may inculcate. *This* sort of decision, made well in advance on the basis of general principles, is not hostage to the need for agents to recognize in the context of judgment that they are biased. Nevertheless, it is a touchstone case of a critical thinking strategy that depends crucially on agents' thinking about the problem. It just enables them to think at arm's length from the situations in which the bias itself will disrupt their capacities to mitigate it.

## V. Closing thoughts

There is a familiar learning model associated with propositional knowledge of the "Paris is the capital city of France" sort. There is also a familiar set of habits of learning and application that enables students (and former students) to apply that knowledge over the longer term of their lives. The problem with IATD is that this sort of knowledge – that biases operate in particular ways, that they can occur in situations like *this one*, the one at hand, and that people are susceptible to them in

contexts like this – does not reliably issue in debiasing behaviours at the point of decision or judgment. A wider view of what counts as successful debiasing indicates a richer class of ways to apply teachable knowledge to the project of debiasing,

Our view, then, is that critical thinking education should encompass practical guidance on how to structure and engage with one’s environment to facilitate good reasoning. This will include teaching how and why to adopt decision-making policies and evidence-gathering practices that do not require the virtuoso ability to rise above invisible and subtle biases. And it will offer learners the opportunity to practice and experiment with this infrastructure creation and reasonable epistemic deference. The intended learning outcomes on our model do include individual learners’ coming explicitly to reason more truth-conducively in specific cases. But they also include, and place great emphasis upon, outcomes that are environmentally mediated, that are implicit and habitual from the individual’s perspective, and that have their main intended effects over the long term and at group levels.

What kind of information, advice, guidance and practice will critical thinking courses of this sort offer? How are these things best taught? These, we think, are among the next big questions in critical thinking education.

## References

- Anderson, C. & E. Sechler. 1986. Effects of explanation and counterexplanation on the development and use of social theories. *Journal of Personality and Social Psychology* 50: 24-34.
- Beaulac, G., & S. Robert. 2011. Théories à processus duaux et théories de l'éducation : le cas de l'enseignement de la pensée critique et de la logique. *Les ateliers de l'éthique* 6.1: 63-77.
- Fischhoff, B. 1982. Debiasing. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press: 422-444.
- Frantz, C. 2006. I AM being fair: The bias blind spot as a stumbling block to seeing both sides. *Basic and Applied Social Psychology* 28.2: 157-167.
- Frantz, C. & R. Janoff-Bulman. 2000. Considering both sides: The limits of perspective-taking. *Basic and Applied Social Psychology* 22: 31-42.
- Gilovich, T. 1991. *How We Know What Isn't So*. New York: The Free Press.
- Goldin, C., & C. Rouse. 2000. Orchestrating impartiality: The impact of "blind" auditions on female musicians. *The American Economic Review* 90.4: 715-741.
- Groarke, L. & C. Tindale. 2004. *Good Reasoning Matters*. Toronto: Oxford University Press.
- Hirt, E. R., & K.D. Markman. 1995. Multiple explanation: A consider-an-alternative strategy for debiasing judgments. *Journal of Personality and Social Psychology* 69: 1069-1086.

- Kenyon, T. forthcoming 2014. False polarization: Debiasing as applied social epistemology. In *Social Dynamics and Collective Rationality, Synthese* special issue, F. Zenker & C. Proietti (eds).
- Kenyon, T. 2008. *Clear Thinking in a Blurry World*. Toronto: Nelson Academic.
- Lilienfeld, S., R. Ammirati & K. Landfield. 2009. Giving debiasing away. *Perspectives on Psychological Science* 4.4: 390-8.
- O'Brien, B. 2009. Prime suspect: An examination of factors that aggravate and counteract confirmation bias in criminal investigations. *Psychology, Public Policy, and Law* 15.4: 315-334.
- Pettigrew, T. F. 1998. Intergroup Contact Theory. *Annual Review of Psychology* 49.1: 65-85.
- Pronin, E. & M. Kugler. 2007. Valuing thoughts, ignoring behavior: The introspection illusion as a source of the bias blind spot. *Journal of Experimental Social Psychology* 43. 4: 565-578.
- Pronin, E., D. Lin, & L. Ross. 2002. The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin* 28: 369-381.
- Pronin, E., C. Puccio, & L. Ross. 2002. Understanding misunderstanding: Social psychological perspectives. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.) *Heuristic and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press. 636-665.
- Reed, R., & A. Sharp (eds.) 1992. *Studies in Philosophy for Children: Harry Stottlemeier's Discovery*. Philadelphia: Temple University Press.
- Ruggiero, V.R. 2004. *The Art of Thinking*. New York: Pearson Longman.

- Sanna, L., S. Stocker & N. Schwarz. 2002. When debiasing backfires: Accessible content and accessibility experiences in debiasing hindsight. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28.3: 497-502.
- Stanovich, K. E., & West, R. F. 2008. On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology* 94.4: 672–695.
- Thaler, R. H., & Sunstein, C. R. 2009. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven: Yale University Press.
- Tetlock, P. 2005. *Expert Political Judgment*. Princeton, NJ: Princeton University Press.
- Thompson, L. 1995. “They saw a negotiation”: Partisanship and involvement. *Journal of Personality and Social Psychology* 68: 839–853.
- Uhlmann, E. & G. Cohen. 2005. Redefining merit to justify discrimination. *Psychological Science* 16.6: 474-480.
- Willingham, D. 2007. Critical thinking: Why is it so hard to teach? *American Educator* 31.2: 8-19.
- Wilson, T. D., & Brekke, N. 1994. Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin* 116.1: 117–142.