

On the Need for Epistemic Enhancement

Democratic Legitimacy and the Enhancement Project

John Danaher

Forthcoming in (2013) 5(1) *Law, Innovation and Technology*

Abstract

Klaming and Vedder (2010) have argued that enhancement technologies that improve the epistemic efficiency of the legal system (“epistemic enhancements”) would benefit the common good. But there are two flaws to Klaming and Vedder’s argument. First, they rely on an under-theorised and under-specified conception of the common good. When theory and specification are supplied, their CGJ for enhancing eyewitness memory and recall becomes significantly less persuasive. And second, although aware of such problems, they fail to give due weight and consideration to the tensions between the individual and common good. Taking these criticisms onboard, this article proposes an alternative, and stronger, CGJ for epistemic enhancements. The argument has two prongs. Drawing from the literature on social epistemology and democratic legitimacy, it is first argued that there are strong grounds for thinking that epistemic enhancements are a desirable way to improve the democratic legitimacy of the legal system. This gives prima facie but not decisive weight to the CGJ. It is then argued that due to the ongoing desire to improve the way in which scientific evidence is managed by the legal system, epistemic enhancement is not merely desirable but perhaps morally necessary. Although this may seem to sustain tensions between individual and common interests, I argue that in reality it reveals a deep constitutive harmony between the individual good and the common good, one that is both significant in its own right and one that should be exploited by proponents of enhancement.

Keywords: Cognitive Enhancement; Epistemic Enhancement; Common Good; Democratic Legitimacy; Epistocracy; Proceduralism; Instrumentalism;

1. Introduction

Buchanan has recently argued for a transformation in the character of the enhancement debate.¹ As he points out, human society has been engaged in an enhancement project for millenia² and very few of the arguments marshalled against the latest biomedical wave of that project are decisive. At best, they offer a set of considerations that we ought to keep in mind as we proceed, however tentatively, with our ongoing enhancement project.

This article develops Buchanan's proposal by looking at the considerations that weigh upon the desirability of biomedical *epistemic enhancements*. That is, enhancements to the ability of humans to acquire knowledge, both theoretical and practical. The article picks up the thread from two pieces by Klaming and Vedder³ which argued (admittedly speculatively) that technologies such transcranial magnetic stimulation might be used to enhance the memory and recall of eyewitnesses presenting evidence before courts and tribunals. The claim was that this would improve the accuracy of legal fact-finding and thereby serve the common good. In presenting this *common good justification* (CGJ) for enhancement, Klaming and Vedder sought to advance the enhancement debate by sidestepping typical objections to enhancement.

But Klaming and Vedder's CGJ for epistemic enhancements is flawed. As others have pointed out, they rely on an under-theorised and under-specified conception of the common good. When the relevant theory and specification is added, their *prima facie* case for epistemic enhancement becomes a good deal less persuasive. Furthermore, although aware of the issue, they fail to accord due weight and consideration to the tensions between the individual and common good that are raised by their proposal. Proper consideration of such tensions is essential if we are to advance the enhancement project.

¹ Buchanan, A. *Beyond Humanity* (Oxford: Oxford University Press, 2011), chapter 2; and *Better than Human?* (Oxford: Oxford University Press, 2012)

² He specifically highlights the examples of literacy and farming. Buchanan n 1 (2011), chapter 2.

³ Klaming and Vedder "Human Enhancement for the Common Good - Using Neurotechnologies to Improve Eyewitness Memory" (2010) 1(3) *AJOB Neuroscience* 22-33; and "Brushing up our memories: Can we use neurotechnologies to improve eyewitness testimony?" (2009) 1(2) *Law, Innovation and Technology* 203-211. The main focus of my discussion is the former piece and the associated peer commentaries.

To that end, this article attempts to revitalise Klaming and Vedder's CGJ for epistemic enhancement. In brief, I argue that Klaming and Vedder reach the right basic conclusion — *i.e.* that epistemic enhancements are desirable and in the interests of the common good — but for the wrong reasons. The argument has two prongs. First, I show how considerations deriving from social epistemology and the literature on democratic legitimacy provide strong *prima facie* reasons for endorsing epistemic enhancement, reasons that draw directly on a conception of the common good. Second, I show how the problems and proposed solutions to the problems associated with scientific evidence move this away from mere desirability into, perhaps, necessity. In doing so, I highlight a deep constitutive relationship between the individual and common good and thereby resolve the tensions highlighted by critics of Klaming and Vedder.

The article proceeds as follows. Section 2 prefaces the main discussion with some terminological clarifications. Section 3 introduces Klaming and Vedder's CGJ for epistemic enhancement and highlights the primary flaws therein. Section 4 presents the two prongs of my argument in favour of epistemic enhancement: the democratic legitimacy prong and the threat of epistocracy prong. Section 5 explains why this argument overcomes the shortcomings associated with Klaming and Vedder's argument. Section 6 concludes by considering how the argument from the preceding sections advances the enhancement project.

2. The Terminology of Enhancement

The enhancement debate is replete with contested and oft-confusing terminology. Particular concerns arise over the distinction between *enhancement* and *treatment/therapy*,⁴ and over the value-laden nature of “enhancement”.⁵ Here, I introduce a new concept --

⁴ See for example Bostrom and Roache “Ethical Issues in Human Enhancement” in Ryburg et al. *New Waves in Applied Ethics* (London: Palgrave Macmillan, 2008); and Savulescu, Sandberg and Kahane “Well-being and Enhancement” in Savulescu, Meulen and Kahane (eds) *Enhancing Human Capacities* (Oxford: Wiley-Blackwell, 2011), for discussions of this issue.

⁵ Harris, J. in *Enhancing Evolution* (Princeton, NJ: Princeton University Press, 2007) and elsewhere seems to suggest that anything that “enhances” must be good, which would of course beg the question in this particular debate. See Brownsword, R. “Regulating Human Enhancement: Things can only get better?” (2009) 1(1) *Law, Innovation and Technology* 125-152 for a response to this argument. See also Pacholczyk & Harris “Dignity and enhancement” in N. J. Palpant & S. C. Dilley (Eds.), *Human Dignity in Bioethics* (London: Routledge, forthcoming) which discusses the problems with the value-laden definition and tries to answer them. Roughly, they argue there is no assumption that the enhancement is *morally* beneficial, but only that it *benefits the individual*.

labelled “epistemic enhancement” -- and argue in its favour. So, given the existing problems with terminology, and the desire to avoid confusions over the interpretation of my argument, it behooves me to offer three clarifications of the terminology adopted in the remainder of this article.

In the first instance, I wish to avoid making value-laden assumptions about nature of “enhancement” by adopting Buchanan’s definition of biomedical enhancement:

Biomedical Enhancement: "A biomedical enhancement is a deliberate intervention, applying biomedical science, which aims to improve an existing capacity that most or all normal human beings typically have, or to create a new capacity, by acting directly on the body or brain".⁶

This definition avoids the charge of automatically assuming enhancement is a good thing by focusing on the intent rather than the effect of the intervention. We can have a genuine moral debate over the propriety of biomedical interventions that are *intended* to improve or add to existing human capacities, without assuming that such improvements are (a) actual or (b) all-things-considered good.

In the second instance, I wish to clarify the different capacities that such biomedical enhancements may target. There are many such targets. Athletes may wish to enhance things like strength, stamina, endurance, and so forth. Others may wish to enhance beauty, or lifespan or height. These, and other cognate forms of enhancement, are excluded from the analysis in this article. The focus here is specifically on enhancements to mental capacities. First, there is a focus on “cognitive enhancement” which, following Bostrom and Sandberg,⁷ may be defined as follows:

Cognitive Enhancement: Any biomedical enhancement targeted at improving the faculties humans use to organise and process information, including (but not necessarily limited to) enhancement of perception, memory, and reasoning.

One potential confusion here is the role that emotional capacities play in cognitive enhancement. There are some theories of emotion that include it within the cognitive

⁶ Buchanan, *Beyond Humanity* (n 1), 23

⁷ “Cognitive Enhancement: Methods, Ethics and Regulatory Challenges” (2009) 15 *Science and Engineering Ethics* 311, p. 311

bracket, but others may exclude it. To ensure that it is not, we can add another category of enhancement, which following Douglas (among others) we can call “moral enhancement”:⁸

Moral Enhancement: Any biomedical enhancement targeted at improving human moral judgment, including (but not necessarily limited to) enhancement of moral emotions (like sympathy and empathy), moral virtues (like courage and generosity) and the impartiality of reasoning processes.

Obviously there is some overlap between the cognitive and moral categories, but also some room for disagreement about the precise constituents of each. For instance, is the emotion of disgust moral and should we encourage it? Evidence suggests⁹ that disgust does indeed play a role in how people reason about moral matters, but there is relatively less agreement about whether it should be encouraged. Some would argue that disgust distorts, rather than improves, moral reasoning. Disagreements of this sort are relevant when considering particular proposals for moral enhancement, but they are less significant here since the goal is to offer a more general endorsement and plan for one part of the enhancement project.

This leads, in the third instance, to the new concept I wish to introduce here. This is the notion of “epistemic enhancement”, which I define as follows:

Epistemic Enhancement: Any biomedical intervention intended to improve or add to the capacities humans use to acquire knowledge, both theoretical and practical/moral.

It is this style of enhancement, and this style alone, that is the focus in the remainder of this article.

3. Vedder and Klaming’s Argument and its Discontents

⁸ Douglas, T. “Moral Enhancement” (2008) 25 *Journal of Applied Philosophy* 228-245; and Pacholczyck, A. “Moral Enhancement: What is it and do we want it?” in Klaming and van den Berg (eds) *Technologies on the Stand Legal and Ethical Questions in Neuroscience and Robotics* (Wolf Legal Publishers, 2011)

⁹ See for instance, Haidt, J. *The Righteous Mind* (London: Penguin Allen Lane, 2012). Haidt proposes a five-part model of moral psychology which includes disgust.

To set-up my proposed CGJ for epistemic enhancement, a contrast with the CGJ offered by Vedder and Klaming will be useful for two reasons. First, it will illustrate the general form and structure that a CGJ for epistemic enhancement can take. And second, by carefully considering the views of Vedder and Klaming's critics, we can see where the CGJ might go wrong. The discussion is divided into three parts. Initially, I look at the proposal made by the authors; then at the argument in favour of that proposal; and finally at the criticisms of that proposal.

3.1 - The Proposal

Vedder and Klaming propose that transcranial magnetic stimulation (TMS) might be used to enhance the memory and recall of eyewitnesses to alleged crimes (or other legally relevant acts). Enhancement of eyewitness recollection would certainly count as an epistemic enhancement. Since memory is an imperfect window into the past, an accurate memory could assist in providing knowledge of past events. This, they argue, would serve the common good since knowledge of the past is essential to the deliberation and decision-making of most legal tribunals.

The proposal to use TMS to enhance eyewitness recollection is admittedly speculative since research on the enhancing effect of TMS is largely in its infancy. But the authors do point to a variety of studies with tantalising results.¹⁰ These include studies suggesting that TMS can be used to disrupt cognitive processes (such as semantic labelling) that increase the incidence of false memories. If this is correct, there may be good grounds for thinking TMS could improve the accuracy of eyewitness memory but it remains to be seen whether this result is robust.

In any event, TMS is just one technology that could lead to epistemic enhancement. Others include the milder form of non-invasive brain stimulation (transcranial direct current stimulation TDCS), which has been found to improve a variety of cognitive processes;¹¹ the invasive form of brain stimulation (deep brain stimulation DBS), which has potential as both a cognitive and emotional enhancer;¹²

¹⁰ Vedder and Klaming (n 3), pp 24-25.

¹¹ Cohen Kadosh, R. et al "The Neuroethics of Non-Invasive Brain Stimulation (2012) 22 *Current Biology* R108-R111, reviews some of the literature on TDCS.

¹² Primarily in the treatment of depression. There are many studies looking at this. A recent example is Holtzheimer, P. et al "Subcallosal Cingulate Deep Brain Stimulation for Treatment Resistant Unipolar

and the more traditional chemical enhancers such as methylphenidate (Ritalin) and modafinil, which have been found to improve concentration and alertness (among other things).¹³ Epistemic enhancement is a complex process and anything that can improve the cognitive and moral capacities alluded to earlier has the potential to act as an epistemic enhancer. So although Vedder and Klaming's test case is TMS, it is important not to exclude the possibility that these other interventions could be justified along similar lines. Furthermore, one should not think that eyewitnesses are the only possible targets of epistemic enhancement. They are merely one among many. Indeed, when I present my own argument in section 4, I will suggest that other participants in the legal system, including judges and jurors, could be potential candidates for epistemic enhancement.

3.2 - *The Argument*

The argument that Vedder and Klaming use to support the proposal is rather opaque. Using a hypothetical example of a fatal hiking accident (person falling off a precipice), the authors suggest that TMS-based enhancement of eyewitnesses could help in confirming or denying the hypothesis that the person was deliberately pushed. The hypothetical is fleshed out in more detail than is possible to provide here, but the basic conclusion is that in this kind of scenario, memory enhancement could assist in confirming or denying legally relevant hypotheses. But to say that is not to say that memory enhancement would be for the common good. Additional assumptions and premises are needed for that, and Vedder and Klaming do not provide them.

Fortunately, one of their critics (Hauskeller) does.¹⁴ He fleshes out the argument in the following manner (I have made one minor modification to Hauskeller's reconstruction in order to avoid compressing one of the key inferences into the conclusion):¹⁵

and Bipolar Depression (2012) 69(2) *Archives of General Psychiatry* 150-158, which was a placebo controlled trial that included a two-year follow-up on efficacy.

¹³ See Repantis, D. et al "Modafinil and Methylphenidate for Neuroenhancement in Healthy Individuals - A Systematic Review" (2010) 62(3) *Pharmacological Research* 186-206, for a review of the evidence on both substances.

¹⁴ Hauskeller, M. "Human Enhancement and the Common Good" (2010) 1(3) *AJOB Neuroscience* 37-39

¹⁵ Hauskeller, *ibid*, p. 38. The change I have made is to divide Hauskeller's conclusion into a principle/premise and a conclusion. I believe this clarifies an additional assumption that underlies the argument.

- (1) Eyewitness testimony “plays an important role in the apprehension, prosecution and adjudication of criminals”¹⁶ because the decisions made by law enforcement officials rely heavily on it.
- (2) Relying on eyewitness testimony can only be justified if it can be trusted, *i.e.* if there is sufficient reason to believe that it is accurate.
- (3) The accuracy of an eyewitness’s testimony depends on the accuracy of their memory, which, however, is notoriously malleable and hence unreliable.
- (4) Therefore, any means of improving the accuracy of memory is desirable with respect to the purpose of apprehending, prosecuting and adjudicating criminals.
- (5) It is in everybody’s interest (except perhaps the criminal’s) that criminals are found out and get convicted, and that innocents are not.
- (6) Therefore, it is beneficial for all of us (except perhaps criminals) for eyewitness memory to be improved (from 4 and 5).

There are a variety of weaknesses in this argument, but I believe it fairly captures the kind of argument that Vedder and Klaming wish to make. The weaknesses shall be discussed in a moment. Before that, however, it is worth briefly noting some of the general features of the argument, and some of the limitations that Vedder and Klaming themselves place upon it.

First, it is worth noting that the argument is neither as specific nor as general as it could be. On the first score, added to this argument might be the specific claim that *TMS has the relevant memory-enhancing effect*, but that more specific claim is neither essential to Vedder and Klaming’s project nor relevant to my concerns in this article. My concern is with the overall case for epistemic enhancement, and the contribution that case makes to the enhancement project. Specific claims about specific technologies can come later. On the second score, a more general argument about the desirability of all sorts of epistemic enhancements, and not just eyewitness memory, would be more in keeping with my concerns here. Nevertheless, the slightly narrower focus on eyewitness

¹⁶ The quote is from Vedder and Klaming (n 3), p. 22

memory, which is the core of Vedder and Klaming's argument, is an instructive case and I shall stick with it for the time being.

Second, it should be noted that, in one respect, Vedder and Klaming's argument is not particularly novel. The suggestion that biomedical enhancement might benefit everyone is not unusual. As others have pointed out,¹⁷ enhancement is not a positional good. That is to say, it is not necessarily something whose value depends on other people not having it. A cognitively enhanced person could create innovations and technological fixes that are beneficial to everyone, even if they are the primary recipients of the benefits. As the saying goes: a rising tide lifts all boats.

Where Vedder and Klaming add some novelty is in expressly directing their argument at an end they feel is in the common interest — *viz.* the apprehension and prosecution of the guilty and the exculpation of the innocent — and pointing out how focusing the argument in this manner avoids some classic pitfalls in the enhancement debate. The first such pitfall is the treatment-enhancement distinction. As they argue, that distinction is really only relevant with respect to the effect of the intervention on the individual, not on the common good. An intervention may simply restore an individual to normal capacity and still be an intervention for the common good. So the relevance of the treatment-enhancement distinction simply falls away in this kind of argument. Furthermore, by focusing on the common good, some of the typical objections to enhancement — those that highlight its self-serving and self-regarding nature — are sidestepped.

Despite these advantages, the argument is limited in a number of ways. As the authors themselves note, the proposal only goes through if (a) the enhancing technology is actually proven to have the desired effect; (b) its side effects (in terms of damage to personal welfare) are minor or non-existent; and (c) concerns about its potential to harm rights to privacy and autonomy can be addressed. The first of these is the most difficult to assess given the relative recency of the technology, but for sake of argument we shall assume the technologies can improve memory (or other epistemically relevant faculties) in the desired manner. As regard the second, Vedder and Klaming discuss the side effects of TMS in their article, noting that they are relatively minor based on what we

¹⁷ For example, see Buchanan, *Beyond Humanity* (n 1), chapter 2 which discusses problems with the framing assumption of assuming that enhancement is primarily a personal good. See also Harris, J. *Enhancing Evolution* (n 5), pp. 28-30; and Bostrom and Sandberg (n 7), pp. 328-329.

now know but that obviously caution is needed. On the final issue, they propose that a system of informed consent could minimise the privacy and autonomy undermining features of this technology. This is something that their critics have targeted and it is to those critics that we now turn.

3.3 - *The Criticisms*

Criticisms of the argument settle around a common theme, one that Vedder and Klaming partly preempt. This theme has to do with the nature of the common good and the clash between it and other personal goods.

The first problem for Vedder and Klaming is that they underspecify their conception of the common good. The common good is a concept that has been subjected to much theorisation over the years.¹⁸ Two major views predominate. The first is that the common good is a *sui generis* or *distinctive* property of the state or polis. In other words, that there is a concept of good that applies to the corporate body known as the state that is distinct from the concept of good that applies to each citizen within the state. According to the second view, the common good is an *aggregative* sum of what is good for each individual citizen. In other words, there is no distinct concept of the common good which is over and above what is good for a person, there is just the aggregation of each individual's good. Charitably, it would seem like Vedder and Klaming are appealing to the second concept since they seem to talk about what is in everybody's interest, not what is in the interest of the state or the polis.

But even if the aggregative concept is what they have in mind, there are problems with its interpretation and application in their argument. This becomes clear if we consider the crucial inference from (4) and (5) to (6) in Hauskeller's reconstruction of their argument. To recap, the inference was the following:

(4) Therefore, any means of improving the accuracy of memory is desirable with respect to the purpose of apprehending, prosecuting and adjudicating criminals.

¹⁸ Particularly in discussions of natural law and political duty. See, for instance, Murphy *Natural Law in Jurisprudence and Politics* (Cambridge University Press, 2006), chapter 3, and Brennan, J. *The Ethics of Voting* (Princeton: Princeton University Press, 2011) chapter 5.

(5) It is in everybody's interest (except perhaps the criminal's) that criminals are found out and get convicted, and that innocents do not.

(6) Therefore, it is beneficial for all of us (except perhaps criminals) for eyewitness memory to be improved (from 4 and 5).

The problem here, as Hauskeller notes,¹⁹ is that just because we all share an interest in one thing (in this instance the successful prosecution of crime) it does not follow that pursuing that interest to the hilt is a good thing. Any aggregative conception of the common good has to factor in possibly competing shared interests and reach some sort of overall balance of those interests. Failing to do so can lead to a *reductio* of the preceding argument. For instance, permanent surveillance and enhanced interrogation techniques²⁰ would no doubt help in preventing and prosecuting crime, but the relentless pursuit of both would not serve the common good. People wish to live in a reasonably secure environment for sure, but they also wish to live in a reasonably free and non-oppressive environment too.²¹ Permanent surveillance and enhanced interrogation would undermine this. These other shared interests have to be taken into consideration so as to reach an all things considered measure of the common good.

This leads nicely to the second main strand of criticism against Vedder and Klaming, namely: that their proposal fails to properly consider damage to personal goods. The clearest expression of this critique is to be found in the response by Chang and Buccafurni.²² As they see it, the CGJ mounted by Vedder and Klaming reveals a deep underlying tension between certain personal interests and the common good. Specifically, they argue that if enhancement²³ is justified on the grounds that it contributes to a commonly shared goal, there is a danger that an individual's sense of self worth will be diminished.

¹⁹ Hauskeller (n 14), p. 38

²⁰ Both examples taken from Hauskeller (n 14), p. 38

²¹ Peter Shiu-Hwa Tsu offers a similar line of criticism when he explores the potentially traumatic impact of memory enhancement on rape victims. Is this something we really wish to encourage just so as to improve the conviction rate? Maybe, but a careful balancing act needs to be undertaken. See Tsu, "Enhancing Eyewitness Memory in a Rape Case" (2010) 1(3) *AJOB Neuroscience* 41-42

²² Chang, P L and Buccafurni, D "Is Invading the Sacred for the Sake of Justice Justified?" (2010) 1(3) *AJOB Neuroscience* 48.

²³ They argue that the term "manipulation" is more appropriate since enhancement, at least as defined by Vedder and Klaming could be harmful to individual welfare. This reveals, once more, the contested value-laden nature of the term "enhancement".

How might this work? The claim is that if it becomes widely accepted that it is legitimate to interfere with a person's cognitive capacities in order to achieve some common goal, then there is a danger that the individual will come to see themselves as a mere cog in a machine; as a mere means to a socially desirable end, not as an end in themselves. This, it is argued, is incompatible with a sense of self-worth:

“[I]ndividual good [is] not simply...the absence of physical or psychological pain but...the presence of respect for cognitive capacities as an intrinsically valuable end in itself. If respect for human cognitive capabilities is treated as an end in itself, then accepting the manipulation of these capacities for the common good is a violation of this respect.”²⁴

Peter makes a similar point, in a different context, but perhaps more forcefully. As she puts it: “people should not simply be seen as patients, who do or do not have well-being, but as agents interested in the autonomous formulation and pursuit of their goals.”²⁵ If they see themselves as mere means to socially desirable ends, they will lose this sense of agency and their sense of self-worth will ebb away.

This looks to be a serious worry. If there is a deep tension between pursuing a shared interest through enhancement and valuing an individual's sense of self-worth, then a CGJ justification for epistemic enhancement would not get off the ground. This is because valuing an individual's self worth and ability to autonomously pursue their own sense of the good is a core commitment of contemporary liberal democracies and anything that undermines this commitment would have to be rejected. Responding to this worry will need to be part of any new positive case for epistemic enhancement.

But, in a way, Vedder and Klaming preempted this kind of concern by admitting that memory enhancement might conflict with rights to privacy and autonomy. Their solution was to adopt a system of informed consent: only those who understood the risks and problems associated with memory enhancement, and consented to those risks and problems, would undergo it. But as a number of critics have argued,²⁶ there are potentially coercive forces associated with enhancement for the common good that might serve to undermine this proposal. For example, in many legal systems, the

²⁴ Chang and Buccafurni (n 22), p. 49.

²⁵ Peter, F. “Pure Epistemic Proceduralism” (2008) 5 *Episteme* 33-55, at 36.

²⁶ Chang and Buccafurni (n 22) pp 49-50; Pierce, R. “Imperfection: Rights, Duties and Obligations” (2010) 1(3) *AJOB.Neuroscience*, 39-41, p. 40; and Lindsay, R. “Don't Forget Memory's Costs” (2010) 1(3) *AJOB.Neuroscience* 35-37.

obstruction of justice is an offence. Arguably, if what is stopping people from helping the police in their investigations is the fact that they have not undergone memory enhancement, they are obstructing justice. The penalties associated with this could have a coercive effect and undermine the informed consent model. Similarly, there is the fact that justice demands a fair trial for all. If this is prevented by the fact that some people choose not to undergo memory enhancement then, even if there is no specific legal requirement for it, there might be a general social pressure to undergo the enhancement. This too could be coercive and undermine the informed consent model.

To sum up, Vedder and Klaming's proposal is subject to at least three lines of criticism, each of which forms a test that a successful CGJ for epistemic enhancement ought to overcome. First, there is the fact that their conception of the common good is underspecified and undertheorised. Assuming they are adopting an aggregative conception of the common good, identifying one shared interest that is served by epistemic enhancement is not enough to show that it serves the common good. Instead, it must be shown that the particular shared interest is balanced with other shared interests so that an all-things-considered determination of the common good is reached. Second, there is a danger that pursuing a shared goal might undermine an individual's sense of self worth. To neutralise this threat it must be shown how epistemic enhancement can contribute to, not detract from, an individual's sense of self-worth. Finally, there is the danger that epistemic enhancement becomes coercive. It must be shown that this is not the case.

4. The Legitimacy Argument for Epistemic Enhancement

In what follows, I develop a CGJ for epistemic enhancement that tries to overcome the problems associated with Vedder and Klaming's. In this section, I set out the basic contours of the argument by initially defending the claim that epistemic enhancement is desirable, before moving on to defend the claim that epistemic enhancement is necessary. In the next section, I show how this argument for epistemic enhancement overcomes the problems associated with Vedder and Klaming's proposal.

The argument develops in three stages. Stage one introduces the set of shared interests that would be served by epistemic enhancements. Stage two shows how

epistemic enhancements can help us to better satisfy those shared interests. Stage three then moves on to show how epistemic enhancement is not merely desirable but actually necessary. In so doing, stage three tries to reveal the deep constitutive relationship between the individual and the common good, at least as it pertains to epistemic enhancement.

4.1 - Liberal Commitments and Legitimacy Conditions

Liberal societies are founded on a core commitment to the freedom and equality of individual citizens. This commitment has been spelled out in different ways. Under the Rawlsian conception, the foundational commitment is embodied by the liberty principle,²⁷ which entitles every citizen to a basic domain of activity that is free from state interference and coercion (*e.g.* freedom of speech, conscience and so forth). Under other conceptions, the foundational commitment is understood in terms of the moral equality of citizens.²⁸ As Gaus explains it, this is the view that no citizen can claim moral authority over another, *i.e.* that they are all equal in terms of their moral authority: I cannot coerce you to follow my moral commands, and you cannot coerce me to follow yours.

What these conceptions share is the notion that freedom from state interference and coercion is, in some sense, the default position, *i.e.* there is no need to morally justify equal freedom, this is the basic right of all persons. Of course, defenders of liberalism — even in the more extreme libertarian form — tend to accept that some coercion is necessary²⁹ in order to avoid conflict and secure the benefits of cooperation. The institutions of state and governance, particularly those of a legal variety, seem to be heavily involved in the creation and maintenance coercive practices: they create coercive rules, and the mechanisms for enforcing those coercive rules. So the key question then becomes how can we justify those institutions in light of the default position?

²⁷ Rawls, J. *A Theory of Justice* (Rev. Edn, Harvard University Press, 1999),

²⁸ See Gaus, G. *Contemporary Theories of Liberalism* (London: Sage Publications, 2003), chapter 8; and *The Order of Public Reason* (Cambridge University Press, 2010), chapter 2

²⁹ In the classic Hobbesian account, the coercive constraints of the state are necessary in order to avoid the state of unending war.

In answering this question, appeal is typically made to *legitimacy conditions*.³⁰ These are conditions which, if satisfied by the decision-making procedures within liberal institutions, render their coercive manifestations (the rules and enforcement mechanisms) legitimate. In other words, they justify the deviation from the default position. Legitimacy conditions break down into two types: *instrumental* and *non-instrumental*. An instrumental legitimacy condition will highlight how a decision-making procedure, or a feature thereof, contributes to a morally acceptable outcome. Hence, it can be justified as a means to a morally desirable end, an end that is in everyone's interest (such as security and beneficial cooperation). A non-instrumental legitimacy condition will highlight how a decision-making procedure, or a feature thereof, is intrinsically morally valuable. Hence, it can be justified in and of itself, without reference to its outputs. This intrinsic value makes it appealing to all citizens, *i.e.* they all have an interest in having procedures of that form. This will typically be because the procedure respects their autonomy and moral agency by giving them a forum for making their preferences known and, if their preferences are not followed, offering reasons to them for the deviation.

It follows from this that there are three different ways to legitimate a decision-making procedure: (i) by appealing only to instrumental legitimacy conditions (*pure instrumentalism*); (ii) by appealing only to non-instrumental legitimacy conditions (*pure proceduralism*); or (iii) by appealing to some combination of both (*mixed legitimation*). We can summarise this in the following principle:

Legitimacy Principle: A coercive rule, practice or institution is justifiable in a liberal society if and only if the procedure through which it is reached satisfies a set of (instrumental/non-instrumental) legitimacy conditions $LC_1 \dots LC_n$.

This allows us to develop the first key premise of my argument in favour of epistemic enhancement. This premise stipulates that it is in everyone's interest to have legitimate decision-making procedures (*i.e.* procedures that follow the legitimacy principle). And this is because legitimate decision-making procedures help secure mutually desirable

³⁰ "Justice conditions" might also be appealed to, but the connection between legitimacy and justice is somewhat disputed. A recent article that explores the connection between justice and legitimacy is: Ceva, E. "Beyond Legitimacy: Can Proceduralism say Anything Relevant about Justice?" (2012) 15 *Critical Review of International Social and Political Philosophy* 183. I agree with Ceva's basic contention that procedures can be a locus of justice that is distinct from outcomes. That could be used to support the argument I advance in the text.

ends, or because they intrinsically respect autonomy and agency (*i.e.* respect individuals as agents not as patients), or because they do some combination of both. In this article, I remain somewhat agnostic as to which approach is correct, hence assuming the mixed account of legitimacy.

4.2 - *The Legitimation of the Trial and Epistemic Enhancement*

The next step of the argument is to show that epistemic enhancement can help governments to be compatible with the requirements of the legitimacy principle. In making this argument, I shall have to proceed carefully, tracing out precisely the mechanisms through which epistemic enhancement can contribute to legitimation. To do this, I will first consider in more detail the form that legitimacy conditions can take in different procedural contexts, specifically in the legislative and trial contexts. I will then try to describe the mechanism through which epistemic enhancement can help satisfy those conditions.

Rawls's four-stage sequence is instructive in this regard.³¹ In setting out a programme for a just society, Rawls specified four procedural contexts, each involved in the construction of rules and principles, and each contributing to the overall structure of the just society. At the first stage — that of the original position — a decision-making procedure is used to locate the basic principles of justice. In Rawls's scheme these are the *liberty principle* (already explained) and the *difference principle* (which specifies how social surpluses and burdens are to be distributed across the populace). At the second stage — that of the constitutional convention — a procedure sets out the framework for respecting basic liberties. At the third stage — that of the legislature — specific rules and laws are introduced for protecting liberties and enforcing a scheme of distribution. And at the fourth stage — that of “everyday life” — the laws of the state are imposed and applied (coercively) on the citizenry. The legal trial is located at this fourth stage as it decides upon the everyday application of the constitutional and legislative rules to the citizenry.

It is not my goal to defend a Rawlsian conception of a just society. But this four-stage schema will be helpful in understanding the form that instrumental and non-

³¹ Rawls, J. *A Theory of Justice* (Rev. Edn, Harvard University Press, 1999), section 31.

instrumental legitimacy conditions can take in the justification of decision-making procedures. The first two stages are too abstract and general to warrant sustained discussion here, and legitimacy conditions arguably play an extremely complex role at these stages.³² So I sidestep a discussion of both by simply assuming that people have a shared commitment to liberty and moral equality. The latter two stages are rather more important and worthy of discussion. Indeed the relationship between legitimacy conditions at both stages plays a significant role in the argument I am making so I shall take some time setting out that relationship here.

Consider initially the legitimacy of legislative procedures, *i.e.* procedures designed to create a specific set of rules for protecting liberties and enforcing a scheme of distribution. Would a dictatorial legislature, in which one person makes all the decisions, be legitimate? How about an elitist legislature, in which an educated elite make all the decisions? Or how about a democratic one, in which all citizens or their chosen representatives make the decisions? One could argue for each, but the generally accepted view is that the democratic legislature would be the most legitimate one,³³ and I'll accept that position here.

The question is “why?”. The answer must come from one of our three justifications listed above (pure instrumentalism, pure proceduralism, or mixed legitimacy). Proponents of epistemic democracy adopt an instrumentalist perspective and argue that democratic procedures are legitimate because they are best at tracking the truth.³⁴ In other words, that democratic procedures — be they based on voting or deliberation — are best at figuring out the normatively correct set of rules.³⁵

Proceduralist defenders of democracy will argue that democratic procedures are

³² Specifically at the first stage there is a worry that the constructive process envisaged by the original position presumes what it wishes to prove, namely: that we ought to be committed to liberty and fair distribution.

³³ It is possible that the others could be too. I'll have more to say about this below, but some of the general arguments are well set out in: Machin, D. “The Irrelevance of Democracy to the Public Justification of Political Authority” (2009) 15 *Res Publica* 103; and also Estlund, D. “Making Truth Safe for Democracy” in Copp, Hampton and Roemer (eds) *The Idea of Democracy* (Cambridge: Cambridge University Press, 1993) and “Why not Epistocracy?” in Reshotko, N (ed) *Desire, Identity and Existence: Essays in Honour of T.M. Penner* (Academic Printing and Publishing, 2003)

³⁴ For a discussion and defence see: Estlund, D. *Democratic Authority* (Princeton, NJ: Princeton University Press, 2008); List and Goodin “Epistemic Democracy: Generalizing the Condorcet Jury Theorem” (2001) 9 *Journal of Political Philosophy* 277 (focusing on the epistemic virtues of voting); and Marti, J.L. “The Epistemic Conception of Deliberative Democracy Defended” in Besson and Marti (eds) *Deliberative Democracy and its Discontents* (London: Ashgate Publishing, 2005) (focusing on the epistemic virtues of deliberation, though not necessarily antagonistic to voting). Gaus, G. *Contemporary Theories of Liberalism* (n 28), chapters 5 and 6 also gives a good overview of the kinds of arguments made by epistemic democrats.

³⁵ Obviously, this presumes cognitivism about moral truths, but this is, as Estlund and others have argued, a relatively weak presumption. Certainly, I will not be challenging the cognitivist assumption here.

intrinsically legitimate because only they can respect moral equality or, more unusually, because they do reach epistemically preferable decisions but that epistemic success can only be assessed by reference to the merits of the procedures, not the outcomes they reach.³⁶ Mixed defences are also, of course, possible³⁷ and, in line with the agnosticism discussed above, will be presumed for present purposes. Mixed defences will take onboard the epistemic claim that democratic procedures are best at tracking the normative truth, but will also argue that there are intrinsically desirable features of democratic procedures, such as their capacity to respect autonomy and agency.

Although my ultimate goal is to show how epistemic enhancements can contribute to the legitimation of trial-based decision-making, I would also submit that they can contribute to the legitimation of democratic legislative procedures. This in itself will constitute a CGJ of epistemic enhancement, but one targeted at a different set of procedures from those targeted by Vedder and Klaming. To see how this is, we must burrow a little deeper into the features of democratic decision-making procedures highlighted by instrumentalists and non-instrumentalists.

Let us first consider the features highlighted by the instrumentalists, those who believe democratic procedures are best at tracking the normative truth. I will not attempt to identify the necessary and sufficient conditions for a procedure's being democratic, but will instead suggest one necessary condition: the *participative* condition. According to this, a procedure is democratic if and only if it allows for nearly³⁸ all those who are affected by its outputs (*i.e.* all citizens) to participate in the process of decision-making. Participation can be achieved via voting or deliberation or both. As it happens, both voting and deliberation are thought to contribute to the epistemic success of democratic procedures and hence to their instrumental legitimation. Voting contributes through the mechanism of the Condorcet Jury Theorem(s). This theorem proves that, provided a certain (increasingly lax) set of conditions holds,³⁹ majority voting is more likely to arrive the correct decision than dictatorial decision-making. Deliberation

³⁶ Peter "Pure Epistemic Proceduralism" (n 25) defends this view.

³⁷ Arguably, Estlund's defence of democracy is of this form. This seems particularly clear in his rejection of epistocracy (rule by an epistemic elite). Although Estlund highlights the truth-tracking potential of democratic procedures, he also argues that rule by an epistemic elite would fail to give citizens the moral reasons they are owed if they are going to be coerced into doing something they would prefer not to do. What is this latter point if not an attempt to highlight the intrinsic virtues of democratic procedures. See Estlund "Why not Epistocracy?" and "Making Truth Safe for Democracy" (both n 33). Of course, there are those who argue that democratic procedures

³⁸ Excluding, as is standard, children who are not of age.

³⁹ The original set of conditions was quite substantial, recent years have seen a relaxation and reduction in the stringency of these conditions. See List and Goodin (n 34) and Marti (n 34) on these points.

contributes to epistemic success by pooling information, allowing for the detection of factual and logical mistakes in reasoning, reducing the manipulation of information and helping to control for irrational preferences.⁴⁰ As Marti argues,⁴¹ voting and deliberation may be mutually reinforcing in terms of their contributions to epistemic success. This would be because deliberative processes increase the chances that the conditions necessary for the application of the Condorcet theorem are met. In a sense then, a system with both deliberation and voting might be epistemically superior when compared to a system with only one or the other.

This brings us to one of the central contentions of this article, namely: that epistemic enhancements could (if they actually work) contribute to the instrumental legitimacy of democratic decision-making procedures. How might this work? Well, very simply. On the deliberative front, epistemic enhancement could allow those who participate to process more information, dampen distorting emotions, remember more facts and so on. This could improve the epistemic efficiency of the deliberative process, allowing it to pool more information and correct for distorting influences. On the voting front, two of the major conditions associated with the Condorcet theorem (and its expansions) are: (i) that the voters be independent of one another and (ii) that they be more likely to vote for the correct outcome than the incorrect one(s).⁴² Again, epistemic enhancement could contribute to the satisfaction of both conditions. It could increase the chances of voter independence by reducing the distorting effect of emotional appeals and, by making individuals more confident (though hopefully not perniciously more confident) in their epistemic abilities, thereby reducing the likelihood of their relying on the views of others. For similar reasons, it could increase their likelihood of voting for the correct answer, *i.e.* by enhancing abilities and reducing distorting effects.

So epistemic enhancement could (if it worked) contribute to the instrumental legitimation of democratic decision-making. Could it do the same for non-instrumental legitimation of democratic decision-making? To answer that we need to consider what it is that makes democratic procedures intrinsically valuable. I will consider the two possibilities highlighted above.

⁴⁰ See Marti (n 34) on these features. Peter “Pure Epistemic Proceduralism” (n 25) also discusses the significant role that deliberative procedures can play in the construction of knowledge.

⁴¹ Marti (n 34)

⁴² As List proves (n 34), this need not mean they have a greater than 0.5 chance of getting the right answer, only that their chance of getting the right answer is greater than their chance of getting any other answer.

The first, following the work of Fabienne Peter,⁴³ suggests that deliberative democratic procedures have intrinsic epistemic merits because they facilitate knowledge construction and respect value pluralism. Peter argues that these epistemic virtues are distinct from those typically endorsed by epistemic democrats because they do not make reference to a procedure-independent standard of correctness. Indeed, part of her argument, drawing upon Deweyan pragmatism about truth, is that there is no independent standard (at least, no *knowable* independent standard) since truth is only determined in light of procedures for arriving at conclusions.⁴⁴ Peter's claim is complex, and its overall plausibility will not be evaluated here, but it essentially breaks down into two sub-claims: the knowledge-construction claim and the respect for value pluralism claim. I shall put the latter to one side for now since it seems to overlap with the respect-oriented views I will discuss momentarily, and focus on the knowledge-construction claim. Here, the notion seems to be that free exchange of ideas and arguments between individuals contributes to the construction of knowledge.⁴⁵ This claim does not seem too far detached from the instrumentalist claim about the epistemic value of deliberation, whatever the underlying metaphysics of truth and knowledge happen to be. The contention is still that the social pooling, testing of ideas, and argument has epistemic merit. Thus, it still seems like epistemic enhancement could have a valuable role to play in increasing epistemic merit. Again, this would be by increasing the capacity for processing and remembering information, and by reducing emotional distortions or biases that prevent the proper consideration and testing of some arguments.

The second possibility is that democratic procedures have merit, not because they contribute to the attainment of knowledge, but because only they can properly respect the autonomy, agency and equality of the citizens. This is a view shared by many theorists. But what does it mean to afford respect to these key properties of individuals? Estlund (and one of his critics, Machin)⁴⁶ see it as matter of providing citizens with the moral reasons they are owed. As noted above, although liberal societies are fundamentally committed to the moral autonomy of citizens, they recognise a need for

⁴³ See in particular "Pure Epistemic Proceduralism" (n 25) and to some extent "The Procedural Epistemic Evaluation of Deliberation" (2012) *Synthese*, DOI: 10.1007/s11229-012-0119-6 (though the later makes a distinct kind of argument in favour of the intrinsic epistemic value of deliberation, one that draws on recent work in the epistemology of disagreement).

⁴⁴ See "Pure Epistemic Proceduralism" (n 25) 41-45 on the Deweyan influence, and section 4 for Peter's own proposed proceduralism.

⁴⁵ *Ibid*, p 43 (on the Deweyan view of knowledge construction) and p. 47-48 (on Longino's view).

⁴⁶ Machin (n 33) is critical of Estlund's claim that only democratic procedures can afford the necessary respect, he is not critical of the claim that such respect needs to be afforded.

some coercive rules and practices. Consequently, people are occasionally going to be forced to do things they would prefer not to. If this is to be legitimate, they must be provided with some moral reasons for doing these things, even if they would not recognise those moral reasons as being decisive. How can this be achieved? Two conditions seem essential: (i) the moral reasons behind the coercive rule or practice must *publicly understandable*; and (ii) the moral reasons must be *reasonably*⁴⁷ (or *qualifiedly*⁴⁸) *acceptable*. The latter condition is complex and contestable, and it is not clear how epistemic enhancement could contribute to its attainment. So I shall not discuss it here. The former is a little more straightforward and epistemic enhancement could play an obvious role in its attainment.

To see this, Machin's example of a legislative schooling policy is useful.⁴⁹ Machin asks us to imagine that the government, via the legislature, introduces a schooling policy that bans private schools and forces parents to send their children to public schools in close proximity to where they live. This means that at least some parents are unable to send their children to the schools they (and the children) would prefer. But the legislature relied on comprehensive reports and costings to decide which schooling policy to introduce. So if they could convey the reasoning underlying their choice to the parents, in a manner that the parents could understand, they would supply them with the moral reasons they are owed. This, however, is extremely difficult to do effectively. As Machin argues, the reports from commissions and expert bodies will be complex and not easily understandable. To overcome this, Machin calls for institutions — actually common in Western democracies — for publicly communicating the relevant information to the parents; and media, think tanks and freedom of information legislation to ensure the information being supplied is fair and unbiased. No doubt this set of information-supplying and scrutinising institutions will always be necessary, but, once again, it seems like epistemic enhancement would assist in rendering the decision-making more perspicuous to the parents. Enhancing their ability to process and recall information, and reducing distorting biases will enable them to better digest and interpret the complex reports underlying the schooling policy. Thus, epistemic enhancement would assist in satisfying the respect condition for democratic legitimacy.

⁴⁷ Machin refers to non-opaque reasons.

⁴⁸ Estlund uses this term in his "Why not Epistocracy?" (n 33) in order to avoid some of the normatively thick concepts associated with the word "reasonable".

⁴⁹ Machin (n 33), pp 114-116

In conclusion then, it looks like epistemic enhancement could improve the legitimacy of legislative decision-making procedures. And since it is in everyone's interest to have legitimate decision-making procedures, epistemic enhancement could, in this respect, serve the common good. As I say, this is a CGJ for epistemic enhancement in its own right, but it targets a different set of procedures than those at issue in Vedder and Klaming's proposal. Could the same basic pattern of reasoning apply to the trial context? In other words, could epistemic enhancement improve the legitimacy of the decisions reached in civil and criminal trials? I turn to this question now.

The trial occupies the fourth stage of Rawls's four-stage sequence. It is concerned with the everyday application of coercive rules and practices to individual citizens. In the criminal context, this involves determining whether a person (or group of persons) committed a crime and therefore ought to be punished. In the civil context, this involves ensuring that agreements are upheld, commitments enforced, compensation is paid and so forth. In both instances, the trial typically has a fact-finding dimension, an interpretive dimension, and an applicative dimension. In reaching decisions about facts and law, courts employ a combination deliberation and voting: evidence and reasoned arguments are presented and discussed by lawyers and judges, and votes are made by juries and panels of judges. In terms of legitimation, both instrumental and non-instrumental justifications of the trial are possible. Instrumentally speaking, the trial could be justified as an "engine for getting at the truth", *i.e.* for ensuring that the guilty are punished and the innocent are not. Non-instrumentally speaking, the trial could be justified by providing a forum in which the individual is given a chance to argue their case (through a representative) and to be presented with the moral reasons they are owed.

Although this description is brief, it highlights how all the essential ingredients for the desirability of epistemic enhancement are present in the trial context. Once again, we have a procedure that is concerned with making the right decisions and respecting the autonomy and agency of the individuals and, as pointed out above, epistemic enhancement could clearly facilitate both. Indeed, its potential facilitation of the former seems to be what implicitly motivated Vedder and Klaming's argument, though they did not tie this into an account of legitimacy conditions. Thus, epistemic enhancement

could improve the legitimacy of the trial and, since legitimation is in the interests of all, we have established the prima facie desirability of epistemic enhancement.

One potential dis-analogy between the legislative and trial contexts is worth addressing before moving on. This concerns the extent of the epistemic enhancement envisaged by the respective arguments. In the legislative context, particularly if we are presuming the desirability of democracy, it seems like epistemic enhancement would be needed for all those who participate in the process, which is potentially everyone in a given society. One could argue that such wide-ranging epistemic enhancement is implausible. But things might be different in the trial context. If we view legitimation in a strictly instrumental fashion then, arguably, a somewhat more selective regime of epistemic enhancement could do the work. In other words, enhancement of just the lawyers, or the judge, or the jury or the witness might be all that is needed to improve accuracy. Enhancement of the litigants or defendants might be largely unnecessary since they typically do not argue and decide the case. Likewise, if we view legitimation in a non-instrumental fashion, then a similarly selective set of epistemic enhancements might be all that is needed. After all, the concern is that the person subjected to the coercive rule (i.e. the litigant or defendant) understand what is going on. Hence it is their epistemic enhancement, not that of the lawyers, judges and juries that really matters.

That said, this claim about the possibility of selective enhancement might be missing some important benefits arising from the epistemic enhancement of all participants. That this is a genuine oversight emerges as we now move on to consider a stronger argument in favour of epistemic enhancement. The argument to this point has established merely the desirability of epistemic enhancement, not its necessity. One could agree with the basic logic of my argument and nevertheless reject the further implication that epistemic enhancement ought to be pursued. One could do so on the grounds that existing procedures are as legitimate as they ever need to be.⁵⁰ Only if we are provided with some reason for thinking that they are not, or that there is serious threat to their legitimacy, will the further implication seem compelling. Providing such a reason is the job of the next sub-section.

4.3 - Epistemic Enhancement as Necessary for Legitimacy

⁵⁰ They could also do so on the grounds that they are not persuaded that the argument overcomes the deficiencies associated

The argument for the necessity of epistemic enhancement derives from Estlund's work on the threat of *epistocracy*. That is: on the potential dangers of handing control of decision-making procedures over to an epistemic elite. Historically, Plato and Mill argued for something along these lines.⁵¹ They did so on instrumentalist grounds, *i.e.* on the grounds that handing control over to an epistemic elite will lead to better decisions and policies. Estlund rejects it by focusing on non-instrumental legitimacy conditions, specifically the need to provide citizens with reasonably acceptable and publicly understandable moral reasons.

In what follows, I explain why an instrumentalist logic seems to support epistocracy. I then outline the two possible solutions to the problem of epistocracy and explain why epistemic enhancement could be a necessary part of both. Finally, I bring the argument out of its purely abstract domain and into the real world, suggesting that epistocracy is not just a possible threat to the legitimacy of the trial but an actual threat. I illustrate this by briefly highlighting the problems posed by the introduction of scientific evidence at trial and specifying how epistemic enhancement could help overcome those problems.

The epistocratic conclusion is derived from a conjunction of three seemingly plausible premises. As follows:⁵²

(1) There are procedure-independent standards of correctness, against which the legitimacy of liberal decision-making procedures ought to be judged.

(Cognitivist Thesis)

(2) In any given society, there will be a group of people with superior epistemic access to these procedure independent standards (**Elitist Thesis**)

(3) If there are people with superior epistemic access to these procedure independent standards, then procedures are more likely to be legitimate if they are given sole or predominant decision-making authority.

⁵¹ Plato *The Republic*; Mill *Considerations on Representative Government* (1861)

⁵² This my own reconstruction, based heavily on that found in Estlund "Making Truth Safe for Democracy" (n 33); *Democratic Authority* (n 34); and Lippert-Rasmussen, K. "Estlund on Epistocracy: A Critique" (2012) *Res Publica*, DOI 10.1007/s11158-012-9179-1

(4) Therefore, in any given society, decision-making procedures are more likely to be legitimate if authority is concentrated in an epistemic elite. (**Authority Thesis**)

If the premises are correct, and if we assume that legitimacy is a morally desirable thing, it looks like we have strong moral reason to support epistocracy. The problem is that this is a deeply unwelcome conclusion, particularly given our liberal commitment to the moral equality of citizens.

The argument can be resisted by rejecting any one of the three premises. However, as Estlund and others point out, this can be difficult to do. The cognitivist thesis is really quite mild, only committed to political truth and justice being independent of what people believe and think, not to some deeply troubling metaphysics of morality. Rejecting it would seem to result in a devastating relativism. The elitist thesis, though perhaps unwelcome, looks to be an uncomfortable truth about most societies. And premise (3) looks like an obvious truth,⁵³ although there may be some ways to resist it by arguing that a larger group would be epistemically superior to a smaller one, even if its members were not individually superior. Finally, the conclusion seems like a natural outgrowth of the cognitivist thesis: if legitimacy is to be judged against these procedure-independent standards then it seems like handing authority to the epistemic elite would increase legitimacy.

So what can we do? Two solutions present themselves. Estlund's is to fall back on a mixed account of legitimacy conditions, one that highlights the need for ensuring that the reasons behind any public decision be acceptable and understandable to the citizenry. His argument is that this could never be ensured by an epistocratic regime as the citizenry would always be left in some doubt as to the actual epistemic skills of the elite. He refers to this as the "Who will know the knowers?" problem of epistocracy and

⁵³ There may be some ways to resist it by arguing that a larger group would be epistemically superior to a smaller one, even if its members were not individually superior. Estlund makes this argument. See, for example, "Why not Epistocracy" (n 33) at p. 56, discussing Aristotle's rejection of a ruling elite. As Estlund goes on to note, Aristotle felt this was too abstract a point since the wiser person will naturally seem to have a superior claim to authority. Likewise, Estlund argues that even if consultation with a larger group were better, the wiser members could be given more votes or time to present their points. This would be in keeping with Mill's suggestion in *Considerations on Representative Government*, which called for more votes for the educated.

suggests that participative democratic procedures are the only way to overcome it.⁵⁴ Those procedures ensure that acceptable reasons are presented to the citizens in an understandable manner, and also give citizens a chance to make their views known.

The difficulty for Estlund's solution is that, as we saw above, epistemic enhancement might be needed if the public understandability condition is to be met. This is particularly so if the decisions being reached are based on complex sets of factual and normative reasoning. If the epistemic abilities of citizens are not enhanced, and if such complexity is present, then one of two things could happen: (i) there could be deep resistance to any decisions made, which might undermine the political structure that was in place; or (ii) there could be undesirable deference to the epistemic elite because too many citizens prefer not to second-guess their judgment. So epistemic enhancement may still be needed to stave off the threat of epistocracy.

The second possible solution to the problem of epistocracy is to accept a basically instrumentalist position, and consider in more depth the uncomfortable truth in premise (2). It would be possible to overcome the threat of epistocracy by raising the level of epistemic abilities across the populace as a whole. Indeed, arguably this is one of the key functions of public⁵⁵ education. And although epistemic enhancements of the sort discussed in section two of this article are not alternatives to public education, they could certainly assist by improving the processing and recall of information. Thus, though it seems likely that epistemic disparities will always exist in some form, the widespread use of epistemic enhancement may help to reduce it and thereby prevent epistocracy.

But this is to speak rather too abstractly. The argument thus far merely suggests that if epistocracy is a real threat, then epistemic enhancement could be a significant part of the solution. But is epistocracy a real threat? It would be too easy to say "yes" and gesture in the direction of the increasing complexity of human knowledge and its

⁵⁴ This conclusion is resisted by some. Machin, for instance, argues that non-democratic procedures could still satisfy the conditions of reasonable acceptability and publicity that Estlund endorses. This disagreement is irrelevant to the argument I am making since my focus is on how those conditions could best be satisfied, not on the argument for democracy. My claim is that epistemic enhancement will (at least sometimes) be necessary to their satisfaction, irrespective of whether the decision-making procedure is, in fact, democratic.

⁵⁵ Here, I am not assuming that "public" means "freely available", just "easily available". In other words, I'm not excluding fee-paying schools and colleges from this calculation.

role in public decision-making procedures. A concrete example, though it will take more time to outline, might persuade readers of the reality of the threat.

The example I have in mind concerns the role of scientific evidence in the trial, the problems associated with it, and some of the solutions to those problems.⁵⁶ Scientific evidence — or rather expert evidence in general — is used to assist a trial in its fact-finding mission (ultimate decisions about the application of the law are left to the judge and jury). For instance, DNA evidence can be used to link a suspect to the scene of a crime; medical evidence can be used to establish injuries in a tort claim; and psychiatric evidence can be used to prove the existence of a “disease of the mind” which may in turn be used to support a defence to a criminal charge.

Such evidence is called upon when the court strays beyond its own areas of expertise and into a domain in which it lacks epistemic competence. This creates two significant problems, both of which ultimately render the threat of epistocracy more real than fictional.

First, there is the problem of *bias*. Experts may be biased in favour of a particular interpretation of the available evidence or the results of their testing protocol for any number of reasons.⁵⁷ This bias distorts the evidence presented and reduces the likelihood of the court reaching a factually correct decision. This is compounded by the second problem, that of *incomprehensibility*. Since courts lack the ability to properly assess the evidence in question, they also lack the ability to properly filter out bias. Hence, we have a kind of trust problem. Courts are not entirely confident in the evidence presented, but because they lack the competence to properly assess it, and because of the general tendency to defer to the authority of the scientific experts, the threat of epistocracy becomes very real indeed. An epistemic elite (in one particular domain) is given the authority to present evidence, often of great significance, to a group of decision-makers who are unable to properly understand it and will tend to defer to it.

⁵⁶ I discuss this in greater depth in Danaher, J. “Blind Expertise and the Problem of Scientific Evidence” (2011) 15 *International Journal of Evidence and Proof* 207

⁵⁷ This is a widely discussed phenomenon. In the field of forensics, Itiel Dror has done trojan work in uncovering systematic biases associated with certain forensic techniques. There are too many works to cite here, but the following is an example: Dror, I., Rosenthal, R. “Meta-analytically quantifying the reliability and biasability of forensic experts.” (2008) 53(4) *Journal of Forensic Sciences* 900-903 He suggests that blinding protocols might be used to overcome some of these problems. I discuss the role of blinding in solving the problem of bias in “Blind Expertise and the Problem of Scientific Evidence” (n 56).

So, in effect, the epistemic elite gain greater control over the decision-making procedures.

Indeed, and somewhat ironically, Vedder and Klaming suggest that epistemic enhancement may compound this problem:

“If TMS is used to enhance the memory of eyewitnesses, law enforcement officials must be able to rely on the expertise and interpretation of findings of experts who use TMS in order to enhance eyewitness memory...[i]n order to rely on an expert, the non-experts must be confident that the expert’s interpretation is in conformity with principles and assumptions that reflect the general opinion in the relevant community of scientists and experts. This is critical when innovative technologies are involved and the common ground between relevant experts [and courts, presumably] and scientists is in the process of being developed.”⁵⁸

What Vedder and Klaming miss, however, is that epistemic enhancements (more broadly conceived than in their example) could play a crucial part in building the necessary “common ground” between the experts and scientists who develop an evidential technique and those required to consider it in a court of law.

We can see this more clearly if we briefly consider one of the more popular solutions to the problems associated with scientific evidence. This is the introduction of a reliability test.⁵⁹ In its original American form, this was referred to as the *Daubert* test, in honour of the case in which it was originally formulated. The test requires the judge to decide whether the scientific evidence is reliable enough to be admitted. To do this, the judge must become acquainted with the indicia of reliable scientific evidence — peer review, error rates, principles and methodologies *etc.* — and consider the evidence presented in light of those indicia. Of course, this requires judges to increase their epistemic competence and would, in order to be effective, require a substantial programme of judicial education. But judges have a lot on their plates already,⁶⁰ and while I wouldn’t wish to denigrate the members of that profession by suggesting they are not up to the job of acquiring the necessary competence, I do wish to suggest that epistemic enhancement could play a valuable role in speeding up the process. If judges

⁵⁸ Vedder and Klaming (n 3), p. 26

⁵⁹ Proposed for introduction in both the UK and Ireland in relatively recent years.

⁶⁰ For one thing, the number of laws they need to familiarise themselves with seems to increase year-on-year. Indeed, this seems to be one of the consequences of having an advanced legal system, particularly in a common law jurisdiction: law *accumulates* over time.

had enhanced information processing and recall abilities, and if they lacked some of the distorting influences of emotion, their path to increased epistemic competence would be less bumpy.

Of course, judicial epistemic enhancement could only be part of the solution. Indeed, creating an epistemically enhanced judiciary, without also creating an epistemically enhanced set of jurors, litigants and lawyers, may simply increase the threat of epistocracy. A non-instrumental understanding of legitimacy conditions could not tolerate that. But even still, as we saw above, epistemic enhancement can be part of the solution to that problem, not the cause of the problem itself. The heavy reliance on scientific evidence is extant. We cannot simply rewind the clock to an epistemically more naive time, nor should we wish to. After all, if we are proponents of instrumentalism or mixed legitimacy, we should want our decision-making to track the truth. But if we want that and wish to overcome the associated threat of epistocracy, epistemic enhancement may be necessary not merely desirable.

4.4 - Summing Up

The preceding discussion has taken a number of important theoretical and practical diversions. But the underlying argumentative thread has been relatively simple. As citizens in a liberal state, we all have an interest in ensuring that our decision-making procedures — particularly the legislative and trial procedures — are legitimate. That is to say, having legitimate decision-making procedures serves the common good. But on any understanding of legitimacy conditions — be it purely instrumental, non-instrumental, or mixed — the epistemic competencies of the citizens are crucial. Thus, epistemic enhancement is at the very least desirable because it increases our chances of satisfying those legitimacy conditions. Furthermore, not only is it desirable, it may actually be necessary in order to fend off the threat of epistocracy (rulership by an epistemic elite). This is contrary to the suggestion by Vedder and Klaming that epistemic enhancement may fuel the fires of epistocracy. But does the argument overcome the other criticisms of their proposal? The next section answers that question.

5. Responding to Criticisms

As we saw in section 3, Vedder and Klaming's CGJ for the use of epistemic enhancement faced several criticisms. While these settled around a common theme — the nature of the common good and its relationship to the individual good — there were three distinct criticisms identified. The first, deriving from the work of Hauskeller, argued that Vedder and Klaming had failed to fully specify the nature of the common good and had thereby neglected tensions between an interest epistemic enhancement and other common interests. Chang and Buccafurni challenged the argument on the grounds that epistemic enhancement of the sort envisaged would undermine the individual's sense of self-worth. And several other critics argued that epistemic enhancement, even if introduced on an informed consent basis, could be coercive and thereby autonomy-undermining. The question now is whether the style of CGJ I have offered for epistemic enhancement can overcome these criticisms.

Before responding to each of them, it is worth clarifying the nature of the common good, and its relationship to the individual good, as it is conceived by my argument. My argument is grounded in a liberal understanding of society. It views the liberty and moral equality of each citizen as paramount. Nevertheless, it recognises that unconstrained liberty would be both impractical and bad. So it is in everyone's interest (*i.e.* it is mutually advantageous) to have some coercive decision-making institutions. But it is also in everyone's interest to make sure that those institutions are legitimate. This requires compliance with the instrumental and/or non-instrumental conditions picked out by the legitimacy principle. It is the satisfaction of these legitimacy conditions that constitutes the (aggregative) common good in my argument. But those legitimacy conditions are valued because they either pick out ends that are in the interest of the individual, or means that respect the individual's morally salient properties. Thus, in the argument, there is an intimate relationship between what is good for the individual and what is good for society as a whole: the one constitutes the other in a non-gappy way. The nature of this relationship is important when assessing the force of the criticisms.

Turning to Hauskeller's criticism first, he took Vedder and Klaming to task for failing to note tensions between the good they identified and other common interests. As he put it, while there may be a common interest in the successful prosecution of crime, this interest cannot be pursued to the hilt. To support this he gave the counterexamples

of perpetual surveillance and enhanced interrogation techniques. While both may assist in the prosecution of crime, their implementation would create an oppressive and undesirable society. The implication then is that relentless pursuit of epistemic enhancement might have a similar effect.

Certainly, no one wishes to live in an oppressive society. But where Hauskeller seems to go wrong is in using as counterexamples to epistemic enhancement, technologies and techniques that create epistemic elites. Surveillance is objectionable because it is centralised and top-down⁶¹ thereby creating information-disparities between government and citizen. Enhanced interrogation is objectionable (in part) because it thrives on a power disparity between the interrogator and the interrogatee. As I have argued, epistemic enhancement, provided it is sufficiently broad-based, could be used to minimise the existence of epistemic elites and hence minimise the kinds of power differentials present in Hauskeller's counterexamples. Thus, far from being a contributing factor to oppression and degradation of other common interests, epistemic enhancement could be used to avoid oppression and preserve our ability to pursue other common interests.

This feeds into an analysis of the second criticism, that of Chang and Buccafurni. They questioned the desirability of TMS memory enhancement on the grounds that it undermined an individual's sense of self worth. The argument was that if the enhancement was being used to ensure more accurate fact-finding, the individual would begin to see themselves as a mere cog in a machine: a mere means to some ultimate end. This would be to view themselves as "patients", *i.e.* recipients of well-being rather than as autonomous agents who create and pursue their own sense of the good life.

Again, this is a serious concern and one that any defence of epistemic enhancement must address. My strategy is to admit that this concern might work to undermine a purely instrumental approach to legitimacy conditions. If it is true that enhanced epistemic abilities are only being used to secure desirable ends, then the individual's sense of self-worth may well be undermined. But as we saw, epistemic enhancement is not only needed for its instrumental benefits, but for its non-

⁶¹ An obvious counterargument is to highlight the potential benefits of a non-centralised, bottom-up approach to information-gathering, *i.e.* sousveillance. David Brin's *The Transparent Society* (New York: Basic Books, 1998) is perhaps the most enthusiastic work about the democracy-sustaining powers of sousveillance.

instrumental benefits too. Epistemic enhancement could be used to sustain the dignity and autonomy of the individual, to ensure that they have a valuable role to play in decision-making procedures, and to ensure that their voice is heard. Thus, far from undermining their sense of self-worth, epistemic enhancement could be used to both sustain and increase it.

Finally, there is the worry about the coercive effects of epistemic enhancement. Coercion is objectionable because it undermines autonomy and responsibility. The concern here is twofold. First, that laws might exist that coerce people into undergoing epistemic enhancement, and second that a general social pressure to undergo epistemic enhancement might be created. This objection looks a little perverse now in light of my argument. The whole point of that argument was to make the case that epistemic enhancement could be used to improve or enhance the legitimacy of coercive decision-making procedures. To turn around and say that epistemic enhancement could itself be illegitimately coercive looks odd. Still, a perversion of this sort could exist: it could be that we undermine the very thing we are trying to achieve. So this objection, like the others, must be taken seriously.

A variety of responses suggest themselves. First, with respect to the coercive effect of the law (e.g. laws on the obstruction of justice), an easy solution would be to create a special dispensation in the case of epistemic enhancement. Provided that the set of epistemic enhancement technologies can be clearly-defined, this shouldn't be too much of a problem.⁶² A second response might be to bite the bullet and argue that some forms of coercion are necessary and that epistemic enhancement is one of those forms. This could be supported with analogies of other, arguably coercive practices, that have similar effects. For example, children are typically coerced into education, which is partly designed to enhance their epistemic abilities, yet this is not deemed objectionable. Indeed epistemic enhancement of the sort envisaged in section one could be rolled out as part of the traditional educational programme.

⁶² One could also argue that some coercion in this domain is acceptable. For example, citizens are *coerced* into submitting themselves for jury duty, and are encouraged to undertake certain decision-enhancing practices when they serve on a jury (e.g. they are asked to avoid reading news stories about the case and to take notes as they hear the evidence). Consequently, adding some degree of biomedical epistemic enhancement to their duties may not be too oppressive a step. Admittedly, depending on the form of enhancement involved, greater risks might be associated with biomedical epistemic enhancements, but that risk can only be assessed on a technology-by-technology basis. I am indebted to Prof. Roger Brownsword for bringing this point to my attention.

A third response is to argue that epistemic enhancement would not be an autonomy-undermining *coercion* at all, but, rather an autonomy-enhancing *offer*.⁶³ Much turns here on how we determine whether or not something is coercive. Following the work of Nozick and Wertheimer,⁶⁴ coercion can be defined as:

Coercion: A coerces B if (i) A makes a *credible threat* to (ii) make B *worse off than they already are* (c) unless B adopts a particular course of action that B would not otherwise have undertaken.

For there to be a genuine instance of coercion, each of these conditions must be met. The paradigmatic instance of coercion might be the highwayman threat of “your money or your life”. In this scenario, the highwayman (i) credibly threatens to (ii) kill the passerby (which makes them worse off than being alive) (iii) unless the passerby hands over money that they would prefer to have kept. By way of contrast, the paradigmatic instance of something that might look coercive but is actually an autonomy-enhancing offer would be a drug company asking a patient to pay money for a life-saving treatment. In this scenario, although the patient has to part with money, they are not being made worse-off than they would otherwise have been (they are going to die; they are being given the chance of living), nor are they clearly being forced to do something that they would otherwise prefer not to do.

The claim then is that epistemic enhancement — at least as it is envisaged by my argument — would not meet the three conditions necessary for coercion. Consider first the desirability argument that I made. This argument was designed to present epistemic enhancement purely as an offer not a threat. It concluded, simply, that if one wished to enhance the legitimacy of decision-making procedures, epistemic enhancement would be a way in which to achieve this. Provided we remove the legal threat associated with obstruction of justice laws, there is nothing coercive about this. The necessity argument was obviously different and appears, on its face, to be somewhat coercive in nature. But I would argue that the necessity argument presented a scenario more akin to that of the patient needing life-saving treatment than the highwayman threatening to kill the passerby. The argument was that epistemic disparities currently exist. And that these

⁶³ I would like to thank Dr. Lesley Haines for drawing this line of argumentation to my attention.

⁶⁴ Nozick, Robert “Coercion,” in Sidney Morgenbesser, Patrick Suppes, and Morton White (eds.) *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, (New York: St. Martin's Press, 1969), pp. 440–472; and Wertheimer, Alan *Coercion* (Princeton, NJ: Princeton University Press, 1987).

disparities increase the threat of epistocracy. Epistemic enhancement was a way of avoiding this outcome. Thus, although the necessity argument looks to have the form “we’d better do this or else”, it is not coercive because it is not making people any worse off. Indeed, quite the opposite: it giving people a chance to escape an undesirable situation in which coercion would become more likely. It is autonomy-enhancing, not undermining. It is an offer, not a coercive threat.

In conclusion then, the CGJ for epistemic enhancement that I offered can potentially overcome the criticisms of Vedder and Klaming’s CGJ. It can do so by adopting a more theoretically sophisticated understanding of the relationship between the individual and the common good. This relationship reveals how, far from increasing the risk of oppression and coercion or undermining an individual’s sense of self-worth, epistemic enhancement can be used to avoid oppression and coercion, and sustain an individual’s sense of self-worth.

6. Conclusion

I conclude by returning to Buchanan’s suggestion, mentioned in the opening paragraph, that the participants in the enhancement debate need to be drawn out of their entrenched pro- and anti- camps – neither of which tend to have categorical arguments in support of their views – and encouraged into a systematic appraisal of the ongoing enhancement project. Drawing a broad historical brush, Buchanan highlights how, from the dawn of tool use, to the birth of agriculture and the spread of literacy, humans have always been enhancing themselves and no doubt always will. The latest biomedical wave of this project is fundamentally no different from the preceding waves, but ethical guidance is always needed to ensure that this project proceeds with an appropriate weighting given to the considerations for and against its particular forms.

Moving beyond the obvious examples of agriculture, literacy and technology, this article has drawn attention to a key aspect of the human enhancement project, namely: the development of democratic systems of governance. Such systems of governance have afforded greater respect to human freedom and moral equality, and have, arguably, ushered in sustained periods of human flourishing and economic development.⁶⁵ What I

⁶⁵ As argued for in Pinker, S. *The Better Angels of our Nature* (London: Penguin Allen Lane, 2011)

have suggested is that a particular range of biomedical enhancements – here called “epistemic enhancements” – can help us to better achieve the democratic ideal and protect that ideal from various threats to its sustainability. And while concerns about the safety of particular technologies will always be present (and should not be overlooked) this democratic aspect of the enhancement project is one that should not be neglected, and ought to be pursued.