

Human Enhancement, Social Solidarity and the Distribution of Responsibility

By John Danaher*

Forthcoming in *Ethical Theory and Moral Practice*

Abstract

This paper tries to clarify, strengthen and respond to two prominent objections to the development and use of human enhancement technologies. Both objections express concerns about the link between enhancement and the drive for hyperagency (i.e. the ability to control and manipulate all aspects of one's agency). The first derives from the work of Sandel and Hauskeller and is concerned with the negative impact of hyperagency on social solidarity. In responding to their objection, I argue that although social solidarity is valuable, there is a danger in overestimating its value and in neglecting some obvious ways in which the enhancement project can be planned so as to avoid its degradation. The second objection, though common to several writers, has been most directly asserted by Saskia Nagel, and is concerned with the impact of hyperagency on the burden and distribution of responsibility. Though this is an intriguing objection, I argue that not enough has been done to explain why such alterations are morally problematic. I try to correct for this flaw before offering a variety of strategies for dealing with the problems raised.

Keywords: Enhancement; Hyperagency; Giftedness Argument; Solidarity; Distribution of Responsibility; Compliance Burdens

* Lecturer in Law, NUI Galway

1. Introduction

The “enhancement project” is the name given to our independent and collective efforts at self-improvement (Buchanan 2011; Hauskeller 2013). The project has been ongoing for some time (Buchanan 2011), and although it has brought with it many benefits, particular objections have been raised to its latest biomedical wave (Sandel 2009; Habermas 2003; Fukuyama 2003; Hauskeller 2013). This article responds to two objections to the enhancement project – here termed “hyperagency objections” – that make significant claims about its effect on the constitutive properties of agency.

The article proceeds in three main phases. Section 2 briefly sketches the general features of a hyperagency objection. Section 3 covers a “solidarity” version of the objection, which has found favour with Michael Sandel and Michael Hauskeller. Both claim that the pursuit of hyperagency will degrade social solidarity because it will undermine our appreciation for the gifted aspects of our lives. I argue that although social solidarity has some value, the mechanism highlighted by Hauskeller and Sandel for its degradation is implausible, and, in any event, the enhancement project can be planned and directed so as to protect solidarity. Section 4 covers a “responsibility” version of the objection, which has been suggested by many but not fully developed by any. The objection holds that in becoming or believing that we will become hyperagents we risk imposing increased burdens or expectations on one another. Again, I try to develop and clarify this objection, in particular by offering a detailed explanation of why it might be wrong to impose additional responsibility burdens on one another. I then argue that although this raises an interesting objection to the enhancement project, it too can be sidestepped.

2. Hyperagency Objections in General

Several philosophers have alluded to the hyperagency problem, but perhaps the clearest illustration comes from a thought experiment by David Owens (2007). The thought experiment takes place in the “Pharmacy of the Future”. In this pharmacy, we can receive medication that alters virtually any of our psychological states. Owens uses the thought experiment to make a point about agency, desire-fulfillment and the enhancement project. On a simple model of agency, an agent is any entity that consists

of: (i) a sensor, which acquires data from the world; (ii) a processor, which combines that data with the agent's internal goals and values; and (iii) an actuator, which allows the agent to act in the world, usually to bring about harmony between the world and its goals/values. Owens's point is that agents typically act within constraints posed by their state of agency. These constraints are technical, cognitive or emotional in nature. Thus, for example, I can't fly to Alpha Centauri, despite my desire to do so, because the technical means for doing so is not available to me yet.

Constraints get in the way of desire-fulfillment, but one thing that technology is good at is removing these kinds of constraints. By allowing us to alter our internal constraints, the Pharmacy of the Future allows us to get more of what we want. This is fine up to point, but it starts to become problematic – or so Owens argues – when we can freely manipulate and enhance any aspect of our agency. His claim is that we need at least some fixed points in order to function as agents at all. If all fixed points are removed, we won't know what it is that we want or should want; we will be landed in a strange, uncertain place. This is the state of hyperagency:

Hyperagency: A state of agency in which nearly every constitutive aspect of agency (i.e. the agent's sensors, processors or actuators) is subject to the agent's will and manipulation.

Several authors seem to be concerned about the pursuit of hyperagency and how the enhancement project is linked to this pursuit. In voicing their concerns, they tend to adopt something like the following argument-template:

Hyperagency Objection Template

(1) The enhancement project will cause us to believe in, pursue or attain a state of *hyperagency*.

(2) If we believe in, pursue, or attain a state of hyperagency, negative consequences $X_1 \dots X_n$ follow.

(3) Any activity that leads to negative consequences $X_1 \dots X_n$ should be forestalled or rethought.

(4) Therefore, the enhancement project should be forestalled or rethought.

Although premise (1) can come in different flavours (specifically, *belief*, *pursuit* and *attainment* flavours), some version of it is common to all hyperagency objections. Where these objections differ is their second premise, specifically in the negative consequence(s) they single out for our consideration and the underlying mechanism of wrong-making. There are two major families of negative consequence discussed by proponents of the objection. The first family keeps the focus on the individual agent, arguing that hyperagency will affect the meaning, significance or value of the individual's life (Danaher 2014). The second family shifts the focus from the individual to the social impacts of hyperagency.

In the remainder of this article, I will focus on two examples from the social family of hyperagency objections. First, I will focus on Sandel and Hauskeller's argument that, by undercutting belief in the giftedness of life, hyperagency will negatively impact on social solidarity. Second, I will develop an argument hinted at by some authors which suggests that increasing opportunities for choice lends itself to increased burdens of responsibility being imposed on other people. Both of these examples fit within the template given above and will be discussed by reference to that template.

3. The Solidarity Version of the Objection

In this section, I examine Sandel and Hauskeller's solidarity version of the hyperagency objection. I do so in three steps. First, I give a general précis of the argument. Second, I try to render more precise the mechanism and the social good – *viz.* solidarity exchanges – that is thought to be undermined by the pursuit of hyperagency. Third, I offer four responses to the objection.

3.1 - What is the Argument?

Sandel's giftedness argument is perhaps the most famous argument in the "hyperagency" family. Hauskeller has refined the argument in light of the leading objections, and it is his refined version of the argument that I focus on here.

The argument starts from a "pursuit"- or "belief"-based grounding premise. Its motivating concern is with the fact that the enhancement project causes us to believe in

or pursue a state of hyperagency, not with our attainment of that state. For the sake of argument, I will grant this grounding premise. This switches focus entirely onto the second premise:

(2) If we believe in or pursue a state of hyperagency, negative consequences X1...Xn follows.

The analytical question then becomes how the variables within this premise (the wrongs), along with the implied mechanism of wrong-making are fleshed out.

We can start with the mechanism of wrong-making. The claim is that belief in or pursuit of hyperagency undermines our appreciation for the gifted aspects of our lives. In his discussion, Hauskeller subjects the notion of “giftedness” to a careful analysis. In one sense, a gifted aspect of one’s life is anything that is *given* or bestowed upon one and treated, more-or-less, as a fixed point within one’s life. For example, the fact that I have two eyes and a mouth is, more or less, a fixed aspect of my life. The fact that I have these features constrains my agency in important ways. The “more-or-less” qualifier is necessary since anything can be modified, destroyed or taken away (ultimately through our deaths), but “more-or-less” these things remain fixed.

In another sense, however, the *gifted* is distinguishable from the given; or, rather, the gifted is a proper subset of the given. A gift, according to Hauskeller, has three core properties: (i) it is good or intended to benefit the receiver; (ii) it cannot be demanded or earned; and (iii) it must be accepted with gratitude. His claim is that certain aspects of our lives — including some of the constitutive properties of agency — are gifts, and ought thus to be treated with appropriate respect, not simply viewed as something to manipulate or alter at our will (as the enhancement project encourages us to do).

There are questions one could raise about the role of giftedness in an atheistic worldview. But Hauskeller tries admirably to defend a secular notion of giftedness and I am willing to concede this defence to him. I accept (*arguendo*) that there are things beyond my control, for which I feel grateful, but which do not (or at least need not) emanate from a divine agent. My interests lie in the link between the degradation of our appreciation for giftedness and the degradation of social solidarity.

This link emanates from the third element in Hauskeller’s analysis of giftedness, *i.e.* with the fact that gifts *must* be accepted with gratitude. Hauskeller argues that being

in receipt of a gift gives rise to a curious and informal obligation to feel grateful (Hauskeller 2011, 64). This obligation, in turn, creates the conditions needed for an important kind of social bonding, namely *solidarity bonding*. Consequently, his claim is that belief in or pursuit of hyperagency will ultimately undermine social solidarity. In other words, Hauskeller fleshes out the second premise of the objection in the following manner:

(2*) If we believe in or pursue a state of hyperagency, we will discourage the appreciation for giftedness, which will block our obligation to feel grateful for the gifted aspects of our lives, which will in turn undermine the conditions needed for social solidarity.

As you can see, there is a two-step mechanism being proposed here (from giftedness to gratitude and from gratitude to solidarity) with the endpoint representing a good-making property of societies. This mechanism is supposedly disrupted by the pursuit of hyperagency.

3.2 - Sharpening the Argument

The mechanism of wrong-making needs further specification, as does the rationale for thinking social solidarity is a great good. Fortunately, these two birds can be killed with the one stone by considering the differences between two kinds of social exchange: (i) economic exchange (or “reciprocal” exchange) and (ii) solidarity exchange (or “altruistic” exchange). The names in brackets are common in the literature, but the types of exchange singled-out by Hauskeller’s argument do not map perfectly onto those concepts. I will stick with the names “economic” and “solidarity” and define the concepts precisely.¹

I’ll start with economic exchange. Imagine Alan wants to buy a car, and Bill wants to sell one. Alan will buy for £1500, and Bill will sell for £1000. After meeting, and haggling, they settle on a price of £1250: Bill gives his car to Alan, and Alan gives the money to Bill. This represents an efficient exchange. Alan gets the car he wants for slightly less than he was willing to pay; Bill gets rid of the car for slightly more than he was willing to accept. It is a win-win. One thing (a car) is directly exchanged for

¹ They do not map on for the reason that solidarity exchanges are not examples of pure altruism as they may involve some expectation of future return. Furthermore, the return expected in an economic exchange is not merely one that is hoped for or expected, it is one that can be legally enforced.

another thing (money). The expectation both parties have of getting something in return for what they offer is built into the very notion of the exchange. So much so, that if you didn't get something in return you would have a legal right to sue for compensation. This last feature of the exchange is particularly important in the present context, and leads to the following definition of an economic exchange (in this definition, and in the next, "A" and "B" could refer to a single agent or to many agents):

Economic Exchange: A engages in an economic exchange with B, if A gives something to B out of a feeling of need or want, believing they have the right to receive something of roughly equivalent value from B in return, and having the ability to enforce that right.

This is to be contrasted with a solidarity exchange, a classic example of which is the voluntary donation of blood. Suppose Alan has a rare blood type, but his life is one of good fortune and good health. He feels grateful for the gift that he has been given, but is aware that there are others less fortunate who might need his rare blood type. He decides to donate his blood. In this manner, he manifests gratitude for his good fortune, and *pays forward* the gift he has been given to others. There may, of course, be the faint worry behind all this that Alan may one day need a blood donation himself, but equally there is the hope that he never will. The blood is given, freely and willingly, with neither the desire nor the demand to receive anything equivalent in return. This leads to the definition of a solidarity exchange:

Solidarity Exchange: A engages in a solidarity exchange with B, if A gives something to B out of a feeling of gratitude, not necessarily directed at B, and without believing or having the right to receive something equivalent in return.

Hauskeller's claim is that solidarity exchanges are a great good-making property of societies. They allow for a type of social bonding to take place that transcends the formal and superficial bonding that takes place in the case of economic exchanges. To quote:

"When I enter a shop, hand over a certain amount of money to the shop attendant and receive some commodity from him in return, we may verbally thank each other, but in fact neither of us will feel particularly grateful to the other [footnote omitted]. And because we don't, the transaction has created no particular social

bond between us, whereas exchanging gifts, by inducing mutual gratitude, does create such a bond.”

(Hauskeller 2011, 65)

As is clear from this, it is the feeling of mutual gratitude that is thought to be the essential lubricant for solidarity exchanges. If people do not have the sense of gratitude for what they have been gifted, they will not have any reason to help others out in a solidarity-enriching manner. This is why the enhancement project is so inimical to solidarity. By encouraging us to view all aspects of our lives, including the traditionally fixed points of our agency, as contingent and manipulable, it undercuts the feeling of gratitude. If someone else is in a worse position than I, it is because they haven't tried hard enough, or manipulated the right things. Our good and bad fortune is something that is (or will be) within our control; it is not something to feel grateful (or ungrateful) about. We thereby lose the motivation for solidarity exchanges.

3.3 - Questioning the Argument

Does the Hauskeller/Sandel version of the hyperagency objection present a serious impediment to the enhancement project, and if so are there ways to work around those impediments? I offer four lines of response.

First, I would cast some doubt on the proposed mechanism of wrong-making. As we saw, the mechanism has two steps: (i) from giftedness to the feeling of gratitude; and (ii) from gratitude to solidarity exchanges. In both cases we must ask: are the causal relationships in question necessary or merely sufficient? In other words, must there be giftedness in order for there to be gratitude, or is it that giftedness is just one really important causal pathway toward gratitude? Similarly, must there be gratitude in order for there to be solidarity, or is gratitude just a really important causal pathway to solidarity? It is hard to say for sure which interpretation of the argument Hauskeller intends, since he does not spell things out in great detail. He does comment that “without gratitude [there is] no reason to help others [in a solidarity compatible way]” (Hauskeller 2011, 67),² suggesting he intends the necessity interpretation (at least in

² I've added the last bit since it seems necessary in order to give a charitable interpretation to Hauskeller. In the original, there is nothing at the end of that sentence, which leads to the odd conclusion that without gratitude there is

relation to the second step), but that seems unfortunate since the necessity interpretation is implausible. This is for two reasons. The first is that there do appear to be grounds upon which gratitude might arise without giftedness. Contrary to Hauskeller's claim, there is no reason why a purely economic exchange couldn't give rise to some feeling of gratitude. In the example given above, Bill may be grateful that he found a willing buyer, and grateful that he got more than he was willing to accept, since the market of potential buyers is large, and their tactics and haggling skills variable. This would seem to be plausible even if he expected something in return, and could legally demand it. Similarly, if I am a rare book collector, I might be very grateful to find a rare book dealer who is willing to engage in an economic exchange with me, knowing as I do that finding such a willing partner is difficult. Nothing in the manipulability of agency would undermine the gratitude found in these cases. Perhaps this economic type of gratitude is less frequent, but it still seems plausible to say that it exists. Similarly, in relation to the second causal step, it may be difficult to prove one way or the other, but it seems likely that empathy, sympathy and trust are the necessary conditions for solidarity exchanges, with gratitude being a very useful catalyst. Alternatively, it could be that none of these things is necessary for solidarity exchanges but they are jointly or independently sufficient.

That suggests the second interpretation (the multiple causal pathways interpretation) is the more plausible and charitable one. And, to be fair, there is something to be said for it. There are studies showing a correlation between gratitude and altruistic pay-it-forward style exchanges (McCulloch et al 2008; Wood et al 2010).³ But there are two interesting observations one can make about such studies. First, there is the suggestion in many of them that gratitude is something that agents can manipulate themselves (e.g. by engaging in a number of behavioural techniques such as keeping a gratitude diary), thereby encouraging more solidarity exchanges, but equally getting a bump in their subjective well-being from doing so. In other words, the studies suggest that gratitude is something we can exercise control over, and which we may have reason to increase, irrespective of solidarity, because of its well-being enhancing effect. Second, these studies don't undermine the suggestion that empathy, sympathy and trust might be alternative causal pathways to solidarity. Indeed, there is another growing body of research on this topic suggesting that it is possible to tweak emotions such as empathy and trust (e.g. through the nasal administration of oxytocin) and thereby

no reason at all to help others. But of course there is such a reason: to get something for oneself. In context, however, I suggest that the sentence does imply the qualifier I have added.

³ Wood et al, in addition to reviewing the evidence offers, practical guidance for increasing the sense of gratitude.

increase solidarity exchanges, though this work is contestable and in its early stages (Savulescu and Persson 2012).

Combined, these two observations point us toward the second, and arguably most important response to Hauskeller's argument. This response holds that, even if he is right, it may still be possible to direct the enhancement project in such a way that solidarity is protected, and perhaps even promoted. This could be done by further enhancing our ability to tweak our sense of gratitude or empathy, and by prioritising research into these types of enhancement. Nothing hinges, incidentally, on the above-suggested mechanisms of solidarity. Though they sound plausible, it may be that other factors are responsible for solidarity, or that solidarity is more easily promoted by other means. All that matters is that it is possible to channel and direct the enhancement project in such a way that it is solidarity-enhancing, not solidarity-undermining.

Two concerns might be raised against this reply to Hauskeller.⁴ In the first instance, one might argue that the reply is a little glib and superficial. It is all too easy to point to hypothetical technological possibilities and use them to respond to critics of enhancement.⁵ Such speculation comes cheap. But how plausible is it that we will be able to successfully manipulate solidarity? Two points seem apposite. First, I concede that there may be technical, empirical and sociological difficulties, but the point I am making is not wholly speculative. There are, as I suggested, active research programmes trying to figure out the ways in which we can manipulate these qualities. I am merely suggesting that we pursue and prioritise these research programmes in order to maintain what Hauskeller deems so valuable. Second, if I am guilty of some empirical speculation in my reply, then so too are the proponents of this style of hyperagency objection. They are trying to map out a general quasi-conceptual objection to the pursuit of hyperagency, based on a particular understanding of the concept of solidarity. I am merely suggesting (using some empirical grounding) that this conceptual understanding could be wrong. There is nothing in the nature of hyperagency that makes it fundamentally incompatible with solidarity or, at the very least, this complaint is not sufficiently credible to completely dissuade us from the pursuit of enhancement.

The second concern is perhaps more significant. It claims that if we manipulate qualities such as empathy and solidarity we risk undermining their authenticity, and their authenticity might be part of what makes them so valuable. One would feel less

⁴ I am indebted to two anonymous reviewers for raising these objections.

⁵ Indeed, I complain about this argumentative style in the enhancement debate in Danaher 2013 [x]

grateful or appreciative (and more suspicious) of someone who has simply manipulated their emotional state in order to facilitate solidarity exchanges. In other words, my suggested reply to Hauskeller could be encouraging us to do something that is ultimately self-defeating of solidarity. This incorporates a fairly standard critique of the enhancement project (Erler 2014; Parens 2014), but maps it onto the particular case of solidarity. This authenticity-objection has some credibility, but there are rebuttals (Douglas 2014).⁶ The authenticity concern can be understood in different ways. It could be based on the notion that the agent in question has bypassed some important moral effort or deliberation in arriving at their present emotional/behavioural state – thus they are not ‘truly’ in that state. Or else it could be based on the notion that their present emotional state is not a reflection of their true authentic selves. The difficulty with these objections is that it is not clear that tweaking empathy or solidarity must defeat solidarity so understood. For starters, it is not clear that the method of enhancement must bypass some important type of moral effort or reasoning. For example, the gratitude diaries that I mentioned above seem to require some effort to write and force the person to reflect more deeply on this morally salient concept. Likewise, the spray of oxytocin doesn’t bypass moral reasoning. It simply heightens awareness of morally salient properties in other people and removes a barrier to moral reasoning. Furthermore, it is not clear that the method of enhancement would bypass the true authentic self. If the tweaking is done at a population level – e.g. by surreptitiously administering oxytocin or by manipulating genetic constitution – then it may, arguably, undermine authenticity (though I have my doubts about the genetic case).⁷ But if the tweaking is left to the individual, then it need not. I may have a higher order desire to be more grateful and to feel more solidarity with my fellow citizens. But I may suffer from certain cognitive and emotional biases that prevent me from realizing that desire. Utilising a method of enhancement may actually bring my behaviour into conformity with my authentic self. In addition to this there is a final conceptual point to be made. The claim that authenticity is what makes solidarity or gratitude so valuable is, at the very least, contestable. It could be that absence of authenticity doesn’t completely undermine its value or that sacrificing this aspect of solidarity allows us to achieve something of considerable value in any event.

⁶ The problem is discussed at much greater length in Douglas 2014. Douglas highlights for different versions of the objection and responds to them at some length.

⁷ The issue is much debated in the literature on free will as to whether such genetic manipulation undermines one’s autonomy. I think there is reason to believe that it does not insofar as one’s self develops after the point of genetic manipulation and one still has (compatibilist) control over one’s reasoning processes, but I accept that this is contested.

This brings me to the third response to Hauskeller's argument. Even if we grant the manipulability of gratitude and solidarity, it is important that we do not become too enamoured with the good-making properties of solidarity in case they distract us from other good-making properties that are easier to achieve. In another part of his critique, Hauskeller laments the fact that enhancement encourages us to make "better" the enemy of "good" (Hauskeller 2011, 71). By constantly encouraging us to think that our lives could be better in the future, enhancement leads us to ignore or downplay what is currently good about them. But his argument on the goods of solidarity exchanges — as things that transcend the economic kind of exchange — also risks downplaying the many positive features of those economic exchanges.

Solidarity is often idealistic and difficult to come by. The pay-it-forward mentality seems to arise in an incentive vacuum. Consequently, it may face an uphill motivational struggle. If we can achieve some of the positive aspects of solidarity exchanges through strictly economic means, we should not shy away from that possibility. For example, if paying people to donate blood increases the blood supply, the fact that doing so might undermine solidarity is not enough reason to dismiss the idea. The positives might outweigh the negatives. This has an important impact on the strength of the preceding objection to Hauskeller's argument. It may turn out that the psychological mechanisms of gratitude and empathy, when tweaked and manipulated to a sufficient degree, no longer allow for pure solidarity exchanges. For instance, it may turn out that people are primarily motivated to engage in such exchanges because of the personal kick they get out of them, not because they bind us closer together. But if some kind of bonding is still possible, and if positive effects flow from that bonding, we should not make the ideal of solidarity the enemy of the good.

This brings us to the fourth and final response to Hauskeller's argument. Despite what I said about channeling the enhancement project toward the promotion of gratitude and solidarity, Hauskeller may insist that the project is simply not compatible with such virtues; that everything I have argued misses the point. He could reason thusly: the tweaking of empathy, or the manipulation of gratitude, will not by themselves drive us to solidarity. The drive toward solidarity depends centrally on the belief that we and others are not directly responsible for our failings; that we are not blameworthy for our misfortune. If that belief is altered, then solidarity will nevertheless be undermined. But this is exactly what the enhancement project does undermine, for it encourages us to see our failings as things we can control and hence as potential grounds for blame. In other words, it encourages us to take an increasingly individualistic and atomistic view: we

are responsible for our individual failings, no one else, and no one deserves our kindness and gratitude simply because they have failed to correct for their own shortcomings.

This is an interesting worry. The impact of enhancement on judgments of responsibility will be discussed in much greater detail in the next section. In doing so, I hope to arrive at a possible solution to this worry. To briefly presage the argument, I am going to suggest that there are ways in which the burdens of responsibility could be distributed (in an enhanced world) that would end up being unfair. But if we are sensitive to this possibility and try to resolve it, we could end up distributing responsibility in a way that enhances social bonding (i.e. enhances the belief that we are all in this together and that we need to work more on one another's behalf). This may still serve to undermine some ideal sense of solidarity, but it still suggests that the enhancement project need not necessarily cause us to become increasingly atomistic and individualistic.

4. The Responsibility Version of the Objection

In this section, I examine a responsibility-based version of the hyperagency objection. I do so in four parts. First, I give a general précis of the argument. Second, I try to strengthen the mechanism that underlies the argument. Third, I clarify the wrong-making properties of that mechanism. And fourth, I offer two potential solutions to the problem.

4.1 - What is the Argument?

The argument I am about to discuss has not been fully articulated by anyone to date, though it is hinted at in some sources. Saskia Nagel (2010) may come closest. Her argument stems from two concerns about hyperagency and the enhancement project. The first concern is the effect that the project can have on the dynamics of choice, and the second concern is the effect it may have on perceptions of blame and responsibility. The two are linked: the effect on the dynamics of choice is what drives the changes in perceptions of blame and responsibility.

The effect on the dynamics of choice can be illustrated by reference to a decision tree. When agents make decisions they usually choose between a variety of possible actions, which can be represented as branches off an original decision point (or “node”).

These choices lead to subsequent decision nodes, which in turn can branch off into more choices, which eventually lead to outcomes. If we adopt a normative decision-theoretical framework, we would typically say that agents should compute the expected utilities of the respective branches along the decision tree in order to decide what to do. This can be an incredibly complicated process as the number of nodes and branches is, potentially, vast. It is pretty clear that human agents don't always perform such calculations when they make decisions. Instead, they rely on heuristics and shortcuts, which allow them to pare the tree down to a manageable size, or skip the calculation altogether and rely on established patterns or habits of behaviour in certain contexts. These peculiarities of human decision-making need not detain us here. What matters for present purposes is the effect that the enhancement project can have on the structure of our decision trees. In essence, it boils down to this: the project can increase, or lead to a perceived or expected increase, in the number of nodes and branches in the decision tree. This tendency to complicate our decision-making by either adding more and more nodes, or encouraging us to believe in or expect more and more nodes, is at the heart of Nagel's hyperagency objection to the enhancement project.

Her claim is that this has a worrying effect when it comes to the mechanisms of blame and responsibility. To use one of her examples:

“If your attention and capacity to concentrate is normal, but does not suffice in very stressful situations in your job, no one can blame you, as long as you try your best. However, if there is modafinil (the drug that helps you stay awake and concentrated) available, and you choose not to take it, you might be blamed for it.”

(Nagel 2010, 114)

The suggestion is that the existence of one cognitive enhancer (in this case modafinil) may alter the dynamics of choice in such a way that what was once upon a time an excusable failure to concentrate is now an inexcusable failure to make use of an enhancement technology.⁸

But is this really problematic? Two problems seem to arise. The first, which Nagel assesses at length in her article, is the effect that a greater feeling of responsibility may

⁸ Others have defended this claim, but have not emphasised the potential badness of this shift in perception, e.g. Santoni di Sio et al (2014)

have on personal well-being. The more choices you have, the more you may feel responsible for your shortcomings and failures, which may lead to a negative self-conception. The second, which Nagel hints at in her article,⁹ is the effect it may have on distributive justice. Put simply, if enhancement technologies help us to minimise risks and improve outcomes, but if some people (through no particular fault of their own) tend to create or have opportunity to create more risks than others, then the burden of risk avoidance may become unequally distributed across society. To the extent that we care, or should care, about distributive justice, this is a concern that needs to be taken seriously.

As I said earlier, although the effect on well-being is worthy of consideration, in this article I am focused on the social effects of hyperagency.¹⁰ As a result, I am going to pursue the second of these concerns. That leads me to flesh out the second premise of this version of the hyperagency objection in the following manner:

(2**) If we believe in, pursue, or attain a state of hyperagency, the burden of risk avoidance may become unequally distributed across members of society.

The mechanism of wrong-making is more opaque than it was when we first filled-in Hauskeller's version of the second premise. In what follows, I attempt to flesh out the mechanism in more detail, and clarify precisely what is wrong with unequal distributions of risk avoidance.

4.2 - Strengthening the Mechanism

Psychological, moral, and legal mechanisms of blame attribution may all contribute to the unequal distribution of risk avoidance. The psychology of blame attribution is complex, and often prone to distorting biases. For instance, the famous Knobe Effect suggests that we are more inclined to attribute an action or outcome to someone when it arises out of "bad" or "negative" intentions than when it arises out of indifferent or good intentions (Knobe 2003; Alfano and Beebe 2012). Still, there are aspects of the psychology which suggest that increased impositions of blame may be

⁹ Nagel (2010) p. 14 at fn 7 says "there are pressing questions concerning distributive justice...that need to be considered carefully. This cannot be done in the present article."

¹⁰ I discuss the effects on personal well-being at length in Danaher (2014).

tied to increased choice. One critical variable in current theories of attribution is the perceived locus of control. When people believe that the locus of control (for preventing or realising some outcome) is outside of themselves or others, they tend not to attribute blame; when they believe the locus of control is internal, they tend to attribute blame. If enhancement technologies lead to the perception of an increasingly internalised locus of control — which they presumably do — then psychological mechanisms of attribution could well lead to increased blame for certain types of failure. What is additionally worrying is that, although people tend to under-attribute blame to themselves, they tend to over-attribute it to others. Thus, while this may limit the impact of enhancement on personal blame, it may actually exacerbate the social impact.

The literature on moral responsibility also provides plausible mechanisms for increased attribution of blame and responsibility. The clearest of these, which is also mirrored in legal doctrine, comes in the shape of the *tracing* condition for responsibility. Since the time of Aristotle, responsibility for an outcome has generally been viewed as a function of two conditions: an *epistemic* condition and a *control* condition. Different accounts of those conditions have been posed over the years but, broadly speaking, they imply that an agent is responsible for an outcome if they knew or reasonably foresaw the outcome, and they had control over that outcome. One problem with this two-condition account of responsibility is its vulnerability to counterexamples like this (modified slightly from Fischer and Tongazzini):

Suppose that Dan gets drunk at Kevin's surprise party and drives home, but his reflexes [and conscious awareness of his activities] are so impaired that he runs over a pedestrian on the way. Whatever sort of control and [epistemic access] is required for moral responsibility, suppose that he didn't have that control [or epistemic access] at the time he ran over the pedestrian. In a straightforward sense, then, he couldn't help it. But is he therefore not responsible?

(Fischer and Tongazzini 2009, 532)

Intuitively, it seems like he is responsible, despite his failure to satisfy the two conditions at the time he ran over the pedestrian. It has been argued that this is because his failure to meet the conditions can be traced back to an earlier moment in time when he had the requisite control over, and/or epistemic access to, the outcome in question (Fischer and Tongazzini 2009; Sher 2009). In this case, that would either be the moment at which he chose to get drunk or chose to drive home. Thus, the classic Aristotelian

conditions of responsibility, however they may be fleshed out, must include tracing-based modifications if they are to effectively capture what it means to be morally responsible. Some dispute the role of tracing in our understanding of moral responsibility (Vargas 2005), but that theoretical dispute does not matter here. All that matters is that tracing seems compelling to many theorists of responsibility and that it could play a role in social decisions about enhancement and responsibility.

I think it is clear that Nagel's concerns about the pursuit of hyperagency map very closely onto tracing based accounts of responsibility. Suppose the enhancement project can and does alter the dynamics of choice in such a way that more decision nodes get added to an agent's decision tree. If these decision nodes represent moments in time at which an agent could have avoided a bad outcome by making use of some enhancement technology, then the tracing condition could well dictate that they are responsible if they fail to make use of it. This is further compounded by the fact that there are existing aspects of legal doctrine that appeal to something like the tracing condition. I'll briefly comment on two of them, one from civil liability and the other from criminal liability.¹¹

In criminal law, an agent is responsible if they satisfy the legally recognised conditions of responsibility (the *elements* of an offence) and cannot avail of a legally recognised defence (justification, excuse or exemption). Examples of excusing or exempting conditions would include things like: loss of control, automatism, insanity, duress and so on. Some of these are partial defences (*i.e.* they reduce a criminal charge but do not allow you to avoid liability completely) and others involve potential restrictions on liberty even if they do provide full exemptions (e.g. insanity). What most of them share is the requirement of some impairment to an agency-relevant capacity such as the capacity for conscious awareness of action, the capacity for rational thought, or the capacity to express one's will in one's actions.

¹¹ The two examples I discuss concern the imposition of liability where the failure to enhance results in harm to another. I focus on these cases because they are ones in which both the moral stakes are high (due to the presence of potential victims) and the potential 'costs' to the person who fails to enhance are high. Nevertheless, in focusing on these two examples, I do not mean to completely exclude the possibility that enhancement could affect other judgments of responsibility in an analogous way. Thus, for example, your failure to cognitively enhance might result in an inability to find work, and society might, consequently, deny you access to social benefits (e.g. job-seekers allowance). The analogy might be that a failure to enhance yourself is similar to a failure to actively seek out employment, and hence the impact on one's access to the social benefit should be similar. The argument I make below about the unfair distribution of compliance burdens could potentially apply to such cases, though perhaps less strongly since the moral stakes seem slightly lower (due to the absence of a potential victim). I am indebted to an anonymous reviewer for suggesting this.

What is interesting about legal practice is that even when these impairments are present, they do not always allow for excuse or exemption. On the contrary: the law often blocks the availability of a defence for reasons that are analogous to those proposed by tracing theorists. To give a legally familiar example, suppose Dan from our earlier tracing case is not a drunk driving home from a party, but an epileptic or a diabetic. Dan manages his condition with medication quite effectively, but when he doesn't he goes into epileptic seizure or hypoglycaemic shock. At that point in time, he no longer has the agency capacities the law requires for responsibility. If he does run over a pedestrian while in such a state is he responsible? Case law suggests that responsibility can be found where the lack of capacity at the time of the offence can be traced back to some prior fault, such as the failure to take his medication (Rumbold and Wasik 2011).¹²

This demonstrates the existence of legal mechanisms whereby the increased choice and increased awareness that arises from the growth of enhancement technologies could allow for increased responsibility. One could counter this by arguing that the cases looked at here do not involve enhancement properly-so-called, but, rather, treatment or therapy. I suspect that there are no principled distinctions between treatment and enhancement in this respect. Nevertheless, I will admit that criminal responsibility is a threshold concept. In other words, one does not necessarily need *enhanced* capacities to be responsible, one just needs "ordinary" capacities. Thus, the impact of the enhancement project on criminal responsibility may be slight. But criminal *liability* — *i.e.* the actual penalty one pays for being criminally responsible — probably is sensitive to degrees of capacity. The more one knows or the more control one has, the greater the price one typically pays for failing to make use that knowledge and control in legally preferable ways. This is borne out when we switch from criminal liability to civil liability.

Consider the mechanisms of liability in relation to professional negligence. In tort law, I am liable to another if I owe them a duty of care and fail to meet the legally prescribed standard of care.¹³ In professional practice, the standard of care is usually fixed by reference to the generally accepted standards of practice within that profession. For example, the medical standard of care in England and Wales, as set down in the

¹² Indeed, Rumbold and Wasik suggest that current criminal law doctrine may be even more punitive than the tracing criterion.

¹³ And if there is a causal link between my breach of duty of care and the injury or harm to the victim. I ignore issues of causation in this discussion.

famous *Bolam* case,¹⁴ holds that a doctor must act “in accordance with a practice accepted as proper by a responsible body of medical men skilled in that particular art”. Where “that particular art” refers to the specific subfield of medical practice the doctor happens to specialise in. It follows from the *Bolam* test that a doctor could be liable if they fail to follow the accepted practices amongst a responsible body of medics.

As others have pointed out, this approach to professional liability in medicine could allow for liability for a failure to use enhancement technologies (Vincent 2011; Goold and Maslen 2014).¹⁵ If an enhancer reliably reduces the risks associated with surgery, it is easy to imagine that the responsible body of medical opinion would shift to the conclusion that performing the surgery without taking the enhancer was not acceptable practice. Consequently, any doctor who failed to take it could be deemed negligent and thus be open to paying a penalty of some kind. Once again there is a plausible legal mechanism whereby increased burdens of risk avoidance are imposed on certain people.

In sum then, the mechanism implicit in Nagel’s argument is a plausible one, as there are psychological, moral and legal ways in which it could be brought about. Is this something we should worry about?

4.3 - Clarifying the wrong-making property

If we grant that the pursuit of hyperagency will *increase* compliance burdens, two questions remain. First, is it true that these increased compliance burdens will become *unequally* distributed? And second, are such distributions morally or ethically problematic?

The first question is easily answered. The mechanisms of blame and attribution target those whose activities have the most potential to bring about socially undesirable ends. For example, the epileptic driver is targeted because their condition presents a greater risk to society than that of a “normal” driver.¹⁶ Likewise, the surgeon is targeted

¹⁴ *Bolam v Friern Hospital Management Committee* [1957] 1 WLR 582

¹⁵ Goold and Maslen argue that currently available cognitive enhancers would not alter professional duties, but they do not rule out the possibility of future enhancement technologies doing so.

¹⁶ I don’t mean to stigmatise or single out epileptic drivers here as this is true for other groups too e.g. young male drivers tend to present greater risks than other drivers. If enhancement technologies can lower the risks they present, they could be targeted in much the same way.

because their activities have greater risks attached to them than those of non-surgeons. Different people have the potential to bring about different kinds of risks, and although there may be benefits associated with risky activities too (a point to which we shall return), assuming there are some pre-existing natural or social inequalities in terms of the risk-potential of different people's activities, it seems obvious that there would be corresponding inequalities in the distribution of the increased compliance burdens.

The second question is slightly more difficult to answer. To the extent that unequal distributions lead to some people being made comparatively worse off (or better off) than others, and to the extent that comparative inequalities are something we prefer to avoid, then there is indeed something morally or ethically problematic about this. But there is reason to be sceptical about the undesirability of comparative inequalities in this instance. One could argue that in the cases discussed above, the imposition of the compliance burden is morally legitimate. For example, in the case of the doctor or the driver the imposition of the burden helps to avoid or minimise harms that could arise, and since harm reduction is a legitimate moral aim, it is entirely appropriate for the law to impose compliance burdens on the doctor or the driver. No one is wronged; society as a whole is better off, not worse. Thus, from the perspective of those planning and guiding the enhancement project, there is nothing to worry about.

Plausible as this seems, I want to suggest that it is mistaken. Relying on some recent work by Andrew Sepielli (2012), I want to argue that comparative inequalities should be matters of concern for those of us planning the enhancement project. Sepielli's work focuses on a simple question: can a law that forbids or requires conduct that *ought to be* forbidden or required, but which has a disproportionate effect on a sub-section of the population, be morally problematic? Sepielli argues that it can, by challenging something he calls *the Offensive Tastes Argument (OTA)*.

According to the OTA, if a law forbids and penalises conduct that it ought to forbid, and this penalty is felt particularly strongly by a sub-section of the population, then this must be because that sub-section of the population has offensive tastes. "Offensive" is defined broadly so as to include all tastes people may have for doing things that they really ought not to be doing. For example, the doctor who refuses to take the cognitive enhancer before the surgery must have a taste for recklessly endangering his patients. He has no right to have that taste and there is no wrong in imposing an increased compliance burden on him.

The problem with this, as Sepielli points out, is that there are at least three kinds of offensive tastes: (i) intrinsically offensive tastes; (ii) extrinsically offensive tastes; and (iii) mere tastes for the offensive.¹⁷ These vary in terms of the blame that can be fairly attached to them. If an agent has intrinsically offensive tastes, then they directly desire to bring about an (objectively) evil or offensive end. There is nothing morally problematic about differentially targeting those with intrinsically offensive tastes. Mere tastes for the offensive are different. As Sepielli defines them, mere tastes for the offensive are tastes that can have an offensive outcome (if satisfied) but the agent is unaware of that possibility and their lack of awareness is not culpable. Sepielli argues that it would be wrong to impose compliance burdens on those with mere tastes for the offensive.

It is the second category of offensive tastes that is the most interesting as it lies somewhere between those two extremes. An agent has an extrinsically offensive taste if: (a) they directly desire something that is permissible or laudable, but (b) the satisfaction of that desire may be offensive, and (c) they are aware of that possibility (or their lack of awareness is blameworthy). Arguably, this would be true of the doctor or the driver in our earlier examples. Presumably, the doctor does not want to recklessly endanger the patient. But, at the same time, she knows that taking the enhancer would improve her chances of success. Thus, she may be aware of the possible bad-making characteristics of attempting surgery without enhancement, but not directly desire those bad-making characteristics. The question is whether imposing a compliance burden on a sub-group of agents with such extrinsically offensive tastes is morally problematic.

The argument is that it can be because of the connection between extrinsically offensive tastes and bad character. Comparative inequality seems justifiable where it is linked to some judgment about an agent's moral character. If the agent is of bad character, then an additional compliance burden is acceptable. The difficulty is that any person can be transformed into a person with (or without) an extrinsically offensive taste, without that transformation seeming to change their moral characters. Consider two people,¹⁸ X who happens to have a taste for *foie gras* and Y who does not. Suppose that there is no humane way to acquire *foie gras*, that both people are aware of this, and that as a result it is legitimately outlawed. In this scenario, X faces a heftier compliance burden than Y due to his innate desire. But if we assume that X's taste is not cultivated,

¹⁷ The discussion of these different tastes can be found in section IV of Sepielli's article.

¹⁸ This example is Sepielli's

nor Y's actively excised or overcome, then this seems unfair. There is nothing intrinsically worse about X's moral character compared to Y.

Why is this? It has to do with the necessary conditions for having an extrinsically offensive taste (EOT), and the modal robustness of character in relation to those conditions. One of the necessary conditions is that the agent *actually has* the taste with extrinsically offensive properties. In the example just given, X has such a taste, Y does not. But adding or removing that taste does not alter our judgment of X's character. We do not think more of him if he lacks the taste, nor do we think less of him for having it. X's moral character is robust across the possible worlds in which he meets that condition and those in which he does not. This holds true for other necessary conditions. Sepielli specifically mentions two: (i) that the fulfillment of the taste has some bad-making properties; and (ii) that the agent actually believes that these properties are bad-making. The second of these is particularly interesting. One would think that if an agent believed that their tastes had bad-making properties, and still acted on them, then this would say something about their character that would justify comparative inequalities in treatment. But there do seem to be cases in which this is not true. Suppose X and Y both have a transmissible disease. X is aware that he has the disease; Y is not. They both choose to come into contact with other people, thereby passing on the disease. On the face of it, X's actions speak to bad moral character, and Y's do not. This might seem to justify differential treatment. But suppose Y is the kind of person that, if he had been aware of the bad-making properties of his disease, would still have chosen to come into contact with other people. In such a case, both X and Y have bad characters and this is consistent across different possible worlds. In some of those worlds, the necessary conditions for EOTs are met, and in some they are not, but because their characters do not change, that in itself doesn't justify unequal treatment.

The important point is that the presence or absence of EOTs does not necessarily say anything to justify differential treatment. That doesn't mean that people who possess EOTs are blameless or that the possession of EOTs is not sometimes expressive of bad character. It just means that the link between EOTs and the justification of differential treatment is loose. In some possible worlds, EOTs speak to moral character; in others, they do not. When the connection is loose, there is reason to worry about unequal compliance burdens.

I argue that this is true in at least some enhancement-related cases. Consider the diabetic driver. In this case, his desire to drive his car is permissible, but has extrinsic

bad-making properties because of his disease. But his having this disease is not blameworthy.¹⁹ When we compare him to another driver who does not have the disease, there is nothing intrinsically worse about their characters. If that's right, then imposing an extra compliance burden on the driver is unfair. It treats him comparatively worse than other drivers for reasons that do not track moral character. The same could be true of the surgeon. Suppose there are two surgeons, X and Y. X wants to do well by his patients, but happens to have a non-rational aversion to taking cognitive enhancers; Y also wants to do well by his patients but is not bothered by taking enhancers. Morally speaking their characters are the same — equally laudable — but if the law imposes a compliance burden on them requiring the use of enhancement, then X is treated comparatively worse than Y, without there being proper character-based justifications for this. (Note: this also works if we compare X, with a member of another profession, Z, who also has an aversion to enhancement, but doesn't face the same kind of compliance burden. One might think the choice of profession does say something significant about the characters, and that this justifies the differential treatment. But this may not be the case if at least some choices of profession are driven by innate or non-cultivated desires analogous to the taste for foie gras. However, the comparison does raise other issues that might affect how problematic the inequality is, as we shall see below).

4.4 - Dealing with the inequalities of compliance burdens

A paradox arises from the preceding analysis. The suggestion seems to be that a compliance burden could itself be legitimate — *i.e.* it could be arrived at following morally and legally acceptable processes of responsibility ascription — but nevertheless raise distributive moral concerns. But how can that be? Surely, the burden is either morally acceptable or not?

The paradox is more apparent than real, based as it is on a one-dimensional approach to both the diagnosis and resolution of moral problems. The assumption underlying the paradox is that the unfairness imposed by the compliance burden can only be resolved by eliminating the compliance burden. But there is no reason to think that this is the only way to resolve the problem. We have to be more creative. Two such creative solutions seem plausible here. The first, suggested by Sepielli in his analysis,

¹⁹ Ignoring some complicating factors arising from Type II diabetes. Assume I talk exclusively of Type I here.

would involve a system of *compensating benefits*. The second, which I add to the mix, involves *spreading the effort burden*. Let's briefly consider both.

A system of compensating benefits would hold that unequal compliance burdens are acceptable, if those who bear them are compensated by some off-setting benefits. These off-setting benefits could come in many different forms. The most obvious and familiar form would be through some sort of direct payment, as is already done through systems of social welfare. Licences to perform actions not generally available to the public might be another kind of off-setting benefit. The benefits need not be funded or channeled through the public purse. For instance, we could argue that in the case of the surgeon being obliged to take cognitive enhancement, there is already a system of compensating benefits in place. The surgeon is well-rewarded for their services, and has a level of respect and job satisfaction that is not available to others. These *might*²⁰ offset the additional compliance burden they face, without imposing any additional burden on public institutions. The case of the epileptic driver is rather different. It is true that being allowed to drive without (additional) fear of prosecution is a kind of benefit that they accrue if they manage their condition with appropriate medications, but in no sense is that a *compensating* benefit. That entitlement only brings them up to the level of other "normal" drivers. Something more would be needed to offset the comparatively unfair compliance burden they must bear. Perhaps, for example, they could receive fuel vouchers or free car repairs.

Although compensating benefits could address the problem of comparative unfairness, there are some problems with them. The main one being that they leave in place the higher standard of behaviour for the individuals in question. If they fail to meet that higher standard, we will punish them accordingly. Thus, if the surgeon does not enhance themselves they (or their insurance fund) may be forced to pay out a significant compensation package to the patient. Similarly, the epileptic driver would be open to criminal liability if they failed to make use of appropriate enhancement technologies. To some extent, this approach makes sense since those individuals look to be best placed to avoid the risks associated with their activities, but there is no reason to think that this is always true, or indeed that the law should always impose the burden on the person best placed to avoid the risk.

²⁰ Whether they do offset the burden is a question that cannot be determined purely in the abstract.

Instead of placing the entirety of the burden on, say, the doctor or the driver in the cases described above, we could place at least some of the burden on other members of society (i.e. we could spread the effort burden). We could do this via new or existing rules of civil or criminal liability, depending on what is appropriate in the context. For example, in the case of the surgeon, other hospital staff (nurses, assisting doctors) might be legally required to remind the surgeon to take the drugs. Or the hospital administration might be required to put in place protocols that ensure that the drugs are taken (e.g. by providing easy access to them, or putting in place warning signs or systems). If these other actors fail in their legal duties, it might take some of the liability away from the surgeon and place it onto these other individuals. Since liability in such cases would be civil, the result could be that the compensation paid to the injured party is paid by actors other than the surgeon. There is nothing particularly unusual in any of this – there are already mechanisms for imposing vicarious liability in civil law systems. My suggestion is merely that these mechanisms be used to spread the effort burden more widely. In the case of the driver, things are slightly trickier. One could imagine legal rules requiring others to warn them or remind them to take their medication, but such a system would be impractical and unlikely to be successful. More plausible, is a legal requirement for car manufacturers to put in place control systems that could prevent people from driving when they have not taken their medication (analogous to existing alcohol interlocks), or which take over when there is some impairment (as is increasingly being made possible through automation).²¹ Failure to put such systems in place (or failure to make them robust enough) may result in liability for the car manufacturers and could also take some of the impact of the burden away from the driver. The type of liability that would be appropriate in such a case is open to debate. One could easily imagine imposing civil liability on car manufactures for such failures (again, legal mechanisms for this already exist). Whether criminal liability is appropriate is another matter, worthy of further consideration. There are other burden-sharing options in these cases too. One might impose requirements on government to properly educate and equip such drivers with the means to avoid the potential risk from their activities. The point here is not to defend a particular model of burden-sharing, merely to highlight that such models are possible and can help mitigate the impact of increased compliance burdens.

²¹ The effect of automation on human activities is beyond the scope of this article, but it could well subvert the need for imposing the compliance burden in the first place. But removing humans from the loop may come with other unwelcome costs.

One interesting scenario – not fully explored to this point – is the case of genetic enhancement and its impact on distributions of responsibility. Suppose that there are means of genetically enhancing one’s offspring which mean that these offspring pose less of a risk to society (e.g. modification of the MAOA gene could reduce risk of violent crime). One might argue, using analogous techniques to those used in the case of the surgeon and driver, that if someone is born and allowed to develop without undergoing such a genetic enhancement, then the parents are responsible for the increased risk. Consequently, the parents should bear a compliance burden that requires them to make use of genetic enhancement. One might argue that this is a particularly strong example of how unfair the distribution of compliance burdens could become.

I agree that this is a particularly strong example of how unfair distributions of compliance burdens might become. I also think that it raises a number of thorny philosophical and political issues (specifically about the obligations of parents toward their offspring and the role of the broader society in influencing or intervening in these decisions). But I think there are some reasons for thinking such a possibility is less likely than the examples discussed previously. For one thing, the link between the parental reluctance/omission and some future wrong is more tenuous than the link between the failure of the surgeon/driver and the resultant harm. Thus, I think it is far less likely that existing mechanisms for imposing legal or moral blame will be leveraged in a way that imposes liabilities on parents. Related to this, I think there are certain goods – parental and reproductive autonomy – that are widely agreed upon which would make us quite reluctant to intervene in reproductive decision-making in this manner. Nevertheless, it is *possible* and if it did happen I think the response should be much the same, *i.e.* we should use some system of compensating benefits or burden-sharing to mitigate the unfairness. This could include, for example, additional benefits to parents who choose to enhance their children in this way, and/or legal requirements on government to put in place systems that provide free access to such enhancements.

This, interestingly, brings us back to Hauskeller’s concerns about social solidarity. As mentioned at the end of the previous section, one concern about the effect of enhancement on the distribution of responsibility is that it would lead us to become more atomistic and individualistic, *i.e.* more inclined to believe that it is up to the individual to correct for their own failings. This might, in turn, erode social solidarity because we will have less reason to think that we are all in this together and need to help each other out. But I hope that in looking at the impact of enhancement on compliance burdens I have given some reason to be more optimistic. In particular, I

hope to have shown that there is reason to be concerned (at a societal level) about the unequal distribution of compliance burdens in light of the increased availability of enhancement technologies. I hope that I have also shown that part of the social solution to that problem is for us to take a less atomistic and individualistic view. If we want the benefits of enhancement (e.g. the benefits of risk reduction) then we need to view this as a less a question of individual responsibility (though that is still important) and more a question of social burden sharing. In short, we need to bind together, not pull apart, if we are to get the most out of the enhancement project. This may not rescue an ideal form of social solidarity, but it may give reason to think the enhancement project could be less corrosive of solidarity than critics seem to suppose.²²

5. Conclusion

The enhancement project is an important one. If pursued wisely, the social and personal benefits could be great: less risk, more choice, more autonomy, greater happiness, and so on. If pursued unwisely, the social and personal costs could be significant. We need to take objections to the project seriously, and plan it wisely. This article has tried to contribute to that effort by focusing on two hyperagency objections to enhancement. The solidarity argument from Sandel and Hauskeller and the responsibility argument, hinted at by the likes of Nagel but more fully developed here. I have argued that neither argument provides overwhelming cause for concern, and that there are ways in which to plan the enhancement project to avoid the potential pitfalls highlighted by these objections.

Acknowledgements – I would like to thank Nicole Vincent, Filippo Santoni de Sio, Philip Robichaud, Nadira Faber and two anonymous reviewers for helpful feedback on earlier drafts of this article.

²² Admittedly, the use of legal means for sharing the effort burden might be more corrosive of solidarity than non-legal means. It would be better if there could be some general shift in cultural attitudes that encourages us to take that view. This may be feasible through some public education and consciousness-raising. That said, the possibility of the law effecting a change in cultural attitudes and beliefs should not be completely discounted.

Bibliography

- Alfano, M and Beebe, J (2012) The Centrality of Belief and Reflection in Knobe Effect Cases: A Unified Account of the Data. *The Monist* 95(2): 264-289
- Buchanan, A (2011) *Beyond Humanity?* OUP, Oxford.
- Douglas, T (2014) Enhancing Moral Conformity and Enhancing Moral Worth. *Neuroethics* 7(1): 75-91
- Danaher, J (2013) The Vice of In-Principlism and the Harmfulness of Love. *American Journal of Bioethics* 13(11): 19-21
- Danaher, J (2014) Hyperagency and the Good Life – Does Extreme Enhancement Threaten Meaning? *Neuroethics* 7(2): 227-242
- Erler, A (2014) Authenticity. In: Bruce Jennings (ed) *Bioethics*, 4th Edition, MacMillan Reference USA, Farmington Hills, MI.
- Fischer, JM and Tognazzini, N (2009) The Truth about Tracing. *Nous*, 43(3): 531 - 556.
- Fukuyama, F (2003) *Our Posthuman Future*. Picador, New York.
- Goold, I and Maslen, H (2014) Must the Surgeon take the Pill? Negligence Duty in the Context of Cognitive Enhancement. *Modern Law Review* 77(1): 60-86
- Habermas, J (2003) *The Future of Human Nature*. Polity Press, Cambridge.
- Hauskeller, M (2011) Human Enhancement and the Giftedness of Life. *Philosophical Papers* 40(1): 55-79
- Hauskeller, M (2013) *Better Humans? Understanding the Enhancement Project*. Acumen, London.
- Knobe, J (2003) Intentional Action in Folk Psychology: An Experimental Investigation. *Philosophical Psychology* 16: 309-324
- McCullogh, Kimeldorf and Cohen (2008) An Adaptation for Altruism? The Social Causes, Social Effects and Social Evolution of Gratitude. *Current Trends in Psychological Science* 17(4): 281-285.
- Nagel, S (2010) Too Much of a Good Thing? *Neuroethics* 3: 109.
- Owens, D (2007) Disenchantment. In Antony, L. (ed) *Philosophers without Gods*. OUP, Oxford.
- Parens, E (2014) *Shaping Our Selves*. OUP, Oxford.
- Rumbold, J and Wasik, M (2011) Diabetic drivers, hypoglycaemic unawareness, and automatism. *Criminal Law Review* 863-872.
- Sandel, M (2009) *The Case Against Perfection*. Belknap Press, Harvard.
- Santoni de Sio, F, Faulmuller, N and Vincent, N (2014) How Cognitive enhancement can change our duties. *Frontiers in Systems Neuroscience* 8: 131.
- Savulescu, J and Persson, I (2012) Moral Enhancement, Freedom and the God Machine. *The Monist* 95(3): 399-421
- Sepielli, A (2012) The Law's 'Majestic Equality'. *Law and Philosophy* doi:10.1007/S10982-012-9165-Y
- Sher, G (2009) *Responsibility without Awareness?* OUP, Oxford.
- Vargas, M (2005) The Trouble with Tracing. *Midwest Studies in Philosophy* 29(1): 269-291
- Vincent, N (2011) The Challenges Posed to Private Law from Emerging Cognitive Enhancement Technologies. In Sam Muller, Stavros Zouridis, Morly Frishman and Laura Kistemaker (Eds). *The Law of the Future and the Future of Law*. Torkel Opsahl Academic EPublisher, Oslo.

- Wood, A, Froh, J.J, Geraghty, AW (2010) Gratitude and Well-being: A Review and Theoretical Integration. *Clinical Psychology Review* 30(7): 890-905.