

In Defence of the Hivemind Society

By John Danaher and Steve Petersen

Forthcoming in *Neuroethics*

Abstract: The idea that humans should abandon their individuality and use technology to bind themselves together into hivemind societies seems both farfetched and frightening – something that is redolent of the worst dystopias from science fiction. In this article, we argue that these common reactions to the ideal of a hivemind society are mistaken. The idea that humans could form hiveminds is sufficiently plausible for its axiological consequences to be taken seriously. Furthermore, far from being a dystopian nightmare, the hivemind society could be desirable and could enable a form of sentient flourishing. Consequently, we should not be so quick to deny it. We provide two arguments in support of this claim – the axiological openness argument and the desirability argument – and then defend it against three major objections.

“We are the Borg. Lower your shields and surrender your ships. We will add your biological and technological distinctiveness to our own. Your culture will adapt to service us. Resistance is futile.”

(Star Trek, *First Contact* 1996)

1. Introduction

Gene Roddenberry’s *Star Trek* presents a utopian vision of the future.¹ The vision is one in which Enlightenment humanism is in the ascendance: individuals are the primary units of moral

¹ For discussions of Star Trek’s utopian pretensions, see Segal, H *Utopias* (Wiley-Blackwell 2012), pp 199 ff; and Saadia, M. *Treconomics* (San Francisco, SF: Pipertext 2016).

concern; their differences and diversities are celebrated;² and the technologically advanced world allows them to flourish. The Borg — an enemy ‘species’ introduced in *Star Trek: The Next Generation* — are the ultimate affront to that ideology. They are a hivemind society, one in which individual differences are suppressed and undermined, and in which the social collective is the unit of moral concern, not the individual. Viewers of the Star Trek universe are supposed to fear the Borg, to see them as the most insidious and terrifying threat to Roddenberry’s utopia. They are not simply an enemy that kills and destroys. They are an enemy that brainwashes and subverts: an enemy that turns you into one of them (or, more correctly, ‘part of them’).³

But should we be so fearful? In this article we argue that we should not. Far from representing an anti-human dystopia, the ideal of a hivemind society is, in fact, one that can be conducive to sentient flourishing. The argument we develop is, of course, speculative and provocative. The goal is not to convert the reader to the hivemind cause but, rather, to encourage them to expand their axiological horizons⁴ and see the hivemind society as one of a number of potentially desirable futures for humanity.

Convincing the reader of this will, no doubt, be an uphill struggle. We make the case in four main stages. First, we try to offer a more precise conceptualisation of a ‘hivemind’ society and thus clarify exactly what it is that we are defending. Second, we offer a preliminary argument in favour of the ideal of the hivemind, one that claims hivemind society would not be as ‘alien’ or ‘inhuman’ as we might first think (or as the example of the Borg might encourage us to think). Third, we present a set of stronger arguments in favour of the hivemind society, that focus on how it can help to bring about certain recognised human goods. Fourth, and finally, we identify and respond to criticisms of this idea.

² Almost: Star Trek is a fictional product of Hollywood and was never ideally progressive when it came to its celebration of human differences and diversities. Homosexual characters, for example, were absent until the series *Star Trek Discovery*.

³ On this distinction, see the Orwell-Huxley exchange on dystopianism. Available at: <http://www.lettersofnote.com/2012/03/1984-v-brave-new-world.html>

⁴ In this sense, the article takes onboard the lesson of Owen Flanagan’s book *The Geography of Morals* (Oxford: OUP 2017), which is to appreciate how large the space of ‘moral possibility’ for humanity actually is.

2. What is a Hivemind Society?

For our purposes, we can distinguish two kinds of hiveminds. In the *broad* sense, a hivemind society is one in which individual agents together form a further agent of some moral significance (the “collective” or, as we prefer, the “hivemind”). As we will see, in this broad sense, we are all plausibly already participating in many such minor hiveminds. In the *narrow* sense, a hivemind society is one in which what otherwise might have been independent agents have no moral value of their own; the agents have “dissolved” into the collective, and are not recoverable as separate entities from the collective. Though it is useful to separate these notions, we are inclined to think there is not in fact a sharp line between these types of hiveminds. The question is really one of degree: to what extent are potentially independent agents merged with a collective, and to what extent has the moral emphasis shifted from the merging agents to the merged agent? The idea that the distinction between a hivemind society and an individualist society is not a sharp or precise one might strike some readers as frustratingly vague or imprecise but we believe that this emphasis on degrees of difference is both metaphysically accurate and, as will become clear, something that can be used to encourage people to take a different perspective on the axiology of hivemind societies: because some of us are, in the broad sense, already tending towards or participating in *hivemind-ish* societies we can at least be open to the possibility of going further in this direction.

This should not be taken to imply that there are no radical differences between hivemind societies and present day societies. At the extreme end of the spectrum there are such radical differences and it is useful to focus on those when considering the argument in favour of hiveminds. For heuristic purposes, then, we will examine the more extreme case of hiveminds in the narrow sense in the remainder of this article. To make the scenario vivid, imagine that a Borg-like collective visits Earth tomorrow. It has been sweeping through the universe “assimilating” countless cultures. Unlike Star Trek’s Borg, however, this collective deigns to recognize our quaint individualism—so instead of simply absorbing us, they offer us each the *choice* to merge with them. We defend here the potential axiological desirability of merging. We hope it is clear, though, that we intend this fanciful scenario as a proxy for a more practical and urgent question. As we will discuss, we think that humans are already drifting in the direction of

ever-more integrated hiveminds, and we think that technology will sooner or later enable humans to join real hiveminds in the narrow, Borg-like sense. Thus the practical moral question is: should we embrace this trend, or resist it? (Assuming, of course, that resistance is *not* futile in this case.)


As we see it, there are two distinct, albeit complementary, axes along which agents might integrate:

Rational Unity: This arises when there is some unity of rationality across the individual human members of a society, i.e. where the group forms a single rational agent. Unity of rationality depends on the unity of the key constituents of rationality which include: beliefs, desires, intentions, deliberative processing, and so forth.

Phenomenological Unity: This arises when there is some unity of phenomenological experience across the individual members of a society, i.e. where in some sense they are seeing, feeling, hearing, touching, or tasting *the same thing*.

Again, both forms of unity come in degrees. It is possible for a society to be more or less rationally unified and more or less phenomenologically unified. The degrees depend on both the length of time the individuals spend in the unified state, as well as the number of properties or experiences they share in the unified state. We can imagine collections of agents with different degrees of either unity, as the graph suggests.

<FIGURE 1 HERE>



It is also worth clarifying what we mean by a ‘society’. For present purposes, we adopt an expansive definition of this concept. A society is simply any grouping or merging of two or more individuals. So when we make the case for the hivemind society we are not necessarily supposing that an entire nation or civilisation of humans joins together into a single collective.⁵ This could happen and if it did it would obviously count as a type of hivemind society, but smaller scale hivemind societies are also possible. It is the degree of integration or unification -- rational or phenomenal -- that matters and not the sheer number of individuals collected together.

Is it currently possible for humans to form hivemind societies? It is not our intention to answer this question. Our intention is to consider whether it is desirable for them to do so. Still,

⁵ Many philosophers will be reminded of Ned Block’s (1978) “Chinese Nation” thought experiment. In that case, though, the mental state of the whole is independent of the mental states of the citizens. Here we consider cases where the mental state of the whole is in some way composed of the mental states of the (former?) individuals.

something has to be said about the overall feasibility of creating hivemind societies, otherwise there is little motivation for the inquiry into their desirability. Consider first the rational unity forms of hivemind society. These would seem to be feasible, since according to many philosophers we already participate in rationally-unified hiveminds (broadly construed). Books such as *Group Agency* (List and Pettit 2011) and *Macro cognition* (Huebner 2014) provide detailed lists of criteria that need to be satisfied in order to ascribe rational unity to a group of humans. To summarize these theories very briefly, it is obviously the case that humans get together to form group agents of a sort (couples, teams, companies, states). Sometimes these group agents achieve quite a high degree of rational unity, sharing and coordinating on goals, beliefs, and intentions in a deeply integrative and interdependent way. When deep integration and interdependency arises, the group rational agency is not simply reducible to the agency of the individual members. Intimate partners, for example, can often sustain rational group agency for extended periods of time. There are conflicts of intention, belief and desire from time to time. But this alone does not undermine the group unity. After all, an individual human being will often experience such conflicts. Still, when it comes to larger-scale rational hiveminds, it seems that the degree of rational unity is often quite tenuous and temporary.⁶ Although some people talk about nations or states as single group agents—a thought perhaps inaugurated by Plato in *The Republic* with his analogy between the ideal city and the human mind, and then later taken up by Thomas Hobbes in *Leviathan* with his vision of the sovereign as an artificial agent made out of many human agents—this frequently stretches credulity given the amount of discord and dissension that is present in the typical state. Nevertheless, modern technologies may enable larger groups to maintain high degrees of rational unity over longer periods of time. The rise of the so-called internet of everything—which involves the incorporation of networked sensors and actuators into everything in the world, including human beings—could enable rational integration and coordination at an unprecedented scale. The historian Yuval Noah Harari argues in his book *Homo Deus* (2016) that this technology enables the deindividuation of humans into a collective.

⁶ Huebner 2014 argues that true largescale hiveminds (he doesn't use this term) are quite rare in practice, but then tempers this by arguing that 'minds' come in degrees and that many group entities approximate or approach hivemind status.

What about phenomenological unity? This is trickier. The idea that two or more humans could share the same phenomenological experience seems alien, possibly even incoherent: phenomenological experiences are sometimes thought of as being necessarily private. But pioneering technologies in direct brain-to-brain communication may render genuine phenomenological sharing possible in the not too distant future. Such an idea goes back at least to Arthur C. Clark (1962), and now is being actively pursued by projects like Elon Musk's *NeuraLink*.⁷ There are already several experiments showing that it is possible to transmit the electrochemical signals from one brain to another. The work by Kevin Warwick (2013 & 2014) and Manuel Nicolelis (2011) provides clear proof of this concept in both human and animal subjects. Furthermore, we know that it is possible to trick a human mind into experiencing things that are happening to someone or something else. The rubber hand illusion, for example, demonstrates that it is possible for one person to 'feel' what is happening to a detached rubber hand (Metzinger 2009, 75-77). An impressive sequence of experiments conducted by a group of European researchers (led by Olaf Blanke, Thomas Metzinger, Mel Slater and Mavi Sanchez-Vives) has taken this principle to an extreme and shown that it is possible to repeatedly trick people into having 'out of body' experiences or to experience things that are happening to a virtual body in a virtual space.⁸ The theory underlying these experiments is that our phenomenological experience of our own body is dependent on the brain constructing a representational model of that body. This model can be disrupted and disoriented by manipulating sensory inputs and feedback loops. This provides a suggestive and encouraging template for anyone who wants to enable a form of 'cloned' phenomenological unity. More speculatively, Kaj Sotala and Harri Valpola, in their article 'Coalescing Minds' (2012) point to various developments in brain-to-brain communication that could enable phenomenological unity, including the possible creation of an 'exocortex'. This would be a prosthetic extension of the brain that can integrate smoothly and seamlessly between two or more brains at the same time, enabling them to function as parts of a single whole. Paul Churchland outlines a similar scenario. He suggests that the corpus callosum, which connects the brain hemispheres in humans,

⁷ For an entertaining summary of the NeuraLink project and its motivations, see the *Wait But Why?* post at <https://waitbutwhy.com/2017/04/neuralink.html>.

⁸ This is the Virtual Embodiment and Robotic Re-Embodiment Project, which produced well over 100 publications over the course of its duration. A full list can be found at <http://www.vereproject.eu/>. A readable summary of some key results is available in Rothman 'Are we already living in virtual reality?', *New Yorker* 3rd February 2018.

already serves the function of a high-bandwidth information channel between two otherwise independent cognitive systems that must *learn* to communicate—and that in principle similar communication could be learned across hemispheres in different humans.

How will such people understand and conceive of other individuals? To this question I can only answer, “In roughly the same fashion that your right hemisphere ‘understands’ and ‘conceives of’ your left hemisphere—intimately and efficiently, but not propositionally!

(Churchland 1988, p. 88)

Such possibilities would obviously enable both greater rational unity as well as greater phenomenological unity.⁹

In short, although it might seem initially farfetched, the creation of a hivemind society has some plausibility. It is possible for humans to be highly rationally integrated with one another, sharing the same intentions, plans and goals and coordinating together to achieve them. Technology is facilitated more extensive forms of rational integration. Furthermore, with current technology, it may soon be possible for humans to be highly phenomenologically integrated. Would high degrees of both forms of integration be desirable? That’s exactly what we consider over the next two sections by developing two arguments in favour of the transition from individualist societies to hivemind societies.

⁹ An interesting edge case is that of conjoined twins. Some historically documented cases of conjoined twins suggest that they can have high degrees of rational unity, and, due to some shared bodily limbs and functions, some degree of phenomenological unity. That said, some of the best documented cases of conjoined twins -- such as the infamous Chang and Eng Bunker twins of the 19th century -- suggest that they can have very different personalities. For a longer discussion see Savulescu and Persson 2016. In addition to this, and as a general point, there are more metaphysical arguments in favour of the possibility of phenomenological unity that are worth considering. Eric Schwitzgebel (2015) argues that, as his title goes, “if materialism is true, the United States is probably conscious”. If he is right—and if materialism is right—then just about all of us are *already* part of phenomenologically unified hiveminds.

3. Why we Should be Open to the Hivemind Society

The first argument can be thought of as a warm-up. It is not intended to prove the desirability of a hivemind society. It is, rather, intended to loosen the imaginative constraints that typically confront people when thinking about the idea of a hivemind society. One of those constraints, felt particularly strongly in Western liberal societies, is that we are currently committed to a kind of individualist moral framework, one in which individuals are the primary and most important locus of moral significance. This is not to say that we are ideologically pure individualists. As we shall point out below, there are anti-individualist, pro-hivemind features of our current moral views, and there are communities and cultures around the world that embrace a more communitarian, anti-individualist outlook. Nevertheless, individualism does seem to be dominant in Western, liberal societies, and the commitment to individualism is often reflected in both moral practices and legal norms of those communities. To free us from the constraints of an individualist outlook, the first argument encourages us to first adopt a principle of axiological openness, *i.e.* a willingness to consider and experiment with different moral possibilities, and then apply that principle to the possibility of a hivemind society.

The principle of axiological openness is a form of historical common sense. We know that moral norms and preferences have shifted many times over the course of human history. If you were a white man living in a European society in the late 1700s it is likely that you would have endorsed slavery and colonialism, thought that women should be denied the legal rights of men, favoured capital punishment and legal torture, and judged homosexuality to be morally abhorrent. While the remnants of these attitudes linger to this day, the radical shift in the moral consensus on these matters since the late 1700s has been remarkable. Very few people openly endorse any of these beliefs today. And this is just a single example of a radical moral shift over time. If we look to other societies and countries, and look deeper into antiquity, the fact of moral change seems undeniable. This should encourage greater humility and uncertainty when it comes to our commitment to our currently favoured set of axiological possibilities (Flanagan 2017; Williams 2015). In other words, the fact of moral variety and change over time and space provides warrant for the principle of axiological openness. To put this in argumentative form:

(1) If moral norms and preferences vary and change dramatically over time (and across cultures), then we should be open to the possibility that our currently favoured set of moral norms and preferences is wrong (and so, by implication, be open to experimenting with different possibilities).

(2) Moral norms and preferences do vary and change dramatically over time (and across cultures).

(3) Therefore, we should be open to the possibility that our currently favoured set of moral norms and preferences is wrong (and be willing to consider or experiment with different possibilities).

Add to this argument the claim that the Western liberal commitment to individualism is just ‘a currently favoured set of moral norms and preferences’ and you get the conclusion that we should be open to axiological alternatives to individualism.

But this attempted application of the principle of axiological openness could face some opposition. While people might be willing to accept *some* principle of axiological openness, they may not accept one that is too open-ended. There is an old adage that you “should keep an open mind but not so open that your brains fall out”. We take it that this means that you need to maintain some standards for rational evaluation in order to properly assess alternative possibilities. This is true when it comes to morality as well. You have to have some evaluative standards and principles against which you can assess the other axiological possibilities. We agree with this. If we abandoned all evaluative principles, then we wouldn’t be able to make the argument in favour of hivemind societies. We would, instead, be landed with a form of radical axiological uncertainty. We feel it is possible to address this concern by appealing to some evaluative principles that have a deep cultural pedigree or seem to be foundational to many moral commitments and worldviews. We can use these to evaluate the desirability of a hivemind future. In other words, we can avoid radical axiological uncertainty by being open to moral change and achieving a “reflective equilibrium” across our axiological and normative beliefs. We

do this, below, by appealing to evaluative principles based on the goods of intimacy, impartiality and meaning.

Against this, however, some might be tempted to argue that the idea of hivemind society is just too deeply alien and different from our own, and that individualism is too deeply ingrained in our current moral worldview, for us to be open to that possibility. In other words, individualism is one of the foundational moral commitments that should shape our evaluation of possible moral outlooks. Tim Christie has argued for something like this in his paper ‘Natural Separateness’ (2009). He does this by presenting a thought experiment that involves a Borg-like society with a high degree of phenomenological unity and asking us to imagine which of our currently favoured moral norms would have to change in such a society. He concludes that “[t]he moral differences between collective morality and human morality would be so substantial that it is hard to imagine” (Christie 2009, 191). We cannot properly assess the desirability of a shift to the hivemind society because it is too different from our current evaluative equilibrium. But is it really that alien? There are, we think, two reasons to remain open to the possibility of a hivemind society.

The first is that the hivemind society may not be that radically different or alien from our current way of life. If you return to the graph from the previous section, and remember the comments made about the possible degrees of unity, you can see that shifting to a hivemind mode of existence need not involve a radical evaluative break from the current status quo. On the graph it is really only the top right hand section (with its high degrees of rational and phenomenological unity) that represents the most radical shift in the space of axiological possibility. The other portions of the graph are much closer to where we currently are and we could imagine transitioning into them slowly, rather than abruptly, by gradually ramping up the degree of rational and phenomenological unity.

In fact, most of modern society seems now to be *choosing* to drift toward the hivemind portions of our graph. (Whether or not this is a *wise* choice we will discuss in the next section.) Already many of us are deeply attached to devices and software that let us connect and share easily with others. Part of what makes connection technology appealing is its ability to facilitate

coordination across teams—whether on the scale of parents texting about who can pick up the sick child, or on the scale of a major corporation distributing its huge tasks through project management software. Such connections can enhance rational unities of various sorts, and we eagerly seek out better and more efficient versions.

Probably the more visceral source of appeal behind connection technology, though, is its ability to help us share *experiences* with others. When we share photos and videos on social media, we are often sharing what we have personally seen, and inviting others to see the same. Even the text we post to social networks is a way to broadcast our thoughts. Of course language is an old, crude, and filtered way for humans to share thoughts—but that seems to be one of its primary functions, and again it is one we are eager to let technology enhance, both in terms of audience breadth and frequency. Broadcasting our private experiences and thoughts across more people more often is plausibly an enhancement to our phenomenological unity.

Now we need simply imagine that connection technology becomes better and easier, greasing the slope toward hiveminds. Suppose for example that a small button implanted on your skull (like those featured in some *Black Mirror* episodes) could allow you to *think* an English sentence like “device, tell my husband I need to get out of this party”, and your husband will receive notification of a direct, private audio message.¹⁰ It could be encrypted in proper ways to ensure privacy, and permit fine-grained choices of who can message you via this method, and when. In effect, at-will telepathy could become a reality. Suppose further that this button is enhanced to the point of an “exocortex” of the type mentioned earlier—one that enables brain-to-brain communication at various *sub*-propositional levels, so that explicitly worded thoughts in a conventional language are no longer required. Direct video and audio streams might be shared from one brain to another, for example, so that grandma could watch with a parent’s eyes while the baby takes first steps. Shared tactile streams would enable a new kind of intimacy. Presumably even emotional responses could be shared, allowing direct and literal empathy between consenting people at any time. Given our current hunger for increased connections, it is hard to see people resisting such technology—at least, if it is made convenient enough. And

¹⁰ Already, for example, there are (ungainly) devices that can read subvocalizations fairly reliably; see <https://www.cnbc.com/2018/04/10/mit-alterego-communicates-with-a-computer-through-subvocalization.html>. The internal audio receiver might work something like the cochlear implants of today.

when one person can read another's thoughts, experience her experiences, and feel her feelings, the lines between minds will be ever more blurry.

Again, it is one thing to say we would be tempted by such increased connectivity, and another to say we should *succumb* to such temptation. But each gradual shift along the path to these different axiological possibilities could be fully warranted and guided by currently favoured evaluative principles, even if this means we end up somewhere radically different by the end. (We develop this thought further in the next section.)

The second reason is that, in the full scope of human history and cultural variation, the centrality of individualism to our axiological worldview is a relatively recent, and somewhat localised development. In his history of European moral philosophy, *The Invention of Autonomy* (1998), JB Schneewind carefully and meticulously documents how the current Western commitment to the moral importance of the individual is something that evolved slowly, over the course of several centuries, not something that was always present. Larry Siedentop, in his book *Inventing the Individual* tells a similar story, describing in particular how Ancient (pre-Socratic) Greek societies did not place the individual at the centre of their morality, instead giving pride of place to the family and the state. These are not unfamiliar worldviews to us even today: various forms of nationalism and communitarianism still strive for our attention. Furthermore, there is at least one major world religion—Buddhism—which is known to favour the view that there is no self in the sense traditionally conceived by Western individualist thought (Flanagan 2017), and which sees the denial of the self as a pathway to Enlightenment. Although there are different interpretations of exactly what the 'No Self' doctrine means, Monima Chadha argues that one leading version of it (the Abhidharma tradition) rejects both the idea of an extended narrative self and a self with any degree of agency and ownership over its actions (Chadha 2017 & 2018). All of these examples suggest that humans are not welded to an individualistic ethos, and so we should be at least open to the possibility of axiological alternatives to individualism.

In conclusion, we believe that the principle of axiological openness is sound and that it encourages us to at least consider and experiment with the desirability of a hivemind society.

4. Why the Hivemind is Desirable

Can a stronger argument be made? There are many axiological experiments we could run. Is there any reason to think the hivemind experiment might be a good one? One that is worth undertaking? We think that there is, and so in this section we argue that the hivemind society is potentially desirable.

‘Desirability’ must be interpreted appropriately. There is a danger that we impose too high an axiological standard on the hivemind society and so rule it out of bounds forever. This could happen if we insist that, in order for it to be desirable, the hivemind society must satisfy important human goods or values and *must not sacrifice or undermine* any other goods or values. This would be an unfair standard. After all, no way of life that is currently deemed desirable meets that standard. Every choice we make involves compromises and tradeoffs. In choosing to be a committed and doting parent, you may have to sacrifice some success in your career. We might prefer it if such compromises could be avoided—and perhaps there are ways to reform social institutions so as to minimise them—but given both finite time and finite resources, some degree of compromise in seeking the good life is inevitable. The same will, necessarily, be true if we favour the hivemind over a more individualistic mode of existence. This does not mean that the hivemind is not desirable. It is enough if we can show that pursuing this ideal satisfies some goods and compensates for significant losses. With that clarification out of the way, we can proceed to the argument we wish to defend, which consists of four distinct sub-arguments in favour of the transition to the hivemind society.

The first sub-argument has to do with the good of intimacy. Achieving an intimate connection with another human being is generally thought to be a core part of the well-lived life. The Aristotelian virtue of friendship, for example, celebrates a bond between friends that involves two people engaging with each other with near-perfect equality and mutuality. Similarly, the ideal of romantic attachment is typically thought to require strong forms of physical and mental mutuality. Our current physical and mental separateness prevents us from achieving near perfect intimacy. There is always some barrier between us. Even in the case of

sexual intimacy—where there is some ‘melting’ of the physical barriers—the mental barriers remain. As the sociologist Ernest Becker once eloquently put it:

“People seem to keep bumping up against each other with their exteriors and falling away from each other...Take even the sexual act — the most intimate merger given to organisms. For most people, even for their entire lives, it is simply a joining of exteriors. The insides melt only in the moment of orgasm, but even this is brief, and a melting is not a communication. It is a physical overcoming of separateness, not a symbolic revelation and justification of one’s interior”

(Becker 1971, 29)

Building the hivemind society, particularly by following the path to high degrees of phenomenological unity, would provide an obvious means for overcoming the mental barriers between us and achieving a more perfected form of intimacy. Furthermore, this is an ideal that could be achieved at different scales and in different forms, depending on both the number of people with whom one forms a hivemind, and the degree of phenomenological sharing it entails.

One might object to this on the grounds that intimacy requires separateness. In other words, one might argue that in order for someone to be intimate with another person, that person and the other person must retain their separate identities. Indeed, it could be that this is what makes intimacy good in the first place. It is like a form of phenomenological ‘edging’: we must be brought to the precipice of complete merger with another but never fall over the edge. The worry might be that the hivemind ideal thrusts us into the abyss. We no longer retain the separateness that makes intimacy such an ecstasy.

Two things can be said in response to this. First, for hiveminds in the broad sense, pursuing the hivemind ideal need not undermine the separateness of individuals. It could itself just be an extreme form of phenomenological edging: humans in hiveminds might still retain some sense of individuality, but this is suppressed or dominated by the sense of merger. Second, even if it does undermine separateness, it is at least disputable as to why this separateness is essential to the

good of intimacy. When we talk about achieving intimacy with another we usually focus on melting the boundaries between us, *i.e.* on trying to peer beneath the exterior mask and get at the person's true identity. Pursuing that ideal to its logical extreme leads us to the narrow hivemind model. There is nothing in the inner logic of intimacy that rules this out.

The second sub-argument has to do with problem-solving and goal achievement. Both of these things are thought to be a core part of the good life (Bradford 2016; Maslen et al 2019). There are different ways to make this case. Some people argue that goal achievement (understood, roughly, as the non-lucky realisation of a desired end) is an intrinsic good, something that is valuable in and of itself. In other words, achieving things through a difficult and non-lucky process, irrespective of what these things might be, is a good and makes life better, all else being equal (Luper 2014 defends something like this). Others, perhaps, more sensibly, argue that achievement is only a good if the goals or ends pursued themselves have value, *i.e.* that the states of affairs brought about through problem-solving and achievement enable us to flourish and thrive. The success of modern, industrialised societies is built on the back of its immense problem-solving capacity, particularly in how we harness energy toward the production of food, clothing, shelter, entertainment and culture (Morris 2015; Tainter 1988). It has long been known that individual humans are not that impressive when it comes to their problem-solving and goal achievement abilities. If you put a chimpanzee and a human in the wild, with no prior training, the chimpanzee will out-live the human every time. Humans depend on the existence of a 'group mind'—a cultural repository of tools, techniques and tricks—to survive. Henrich (2015) documents this at length giving numerous examples of how important 'collective intelligence' is to human flourishing. What's more, it seems like collective intelligence is becoming more and more important to our flourishing in the modern era. Numerous studies suggest that our capacity for research and innovation is subject to diminishing marginal returns (Bowles et al 2018; Nielsen and Collison 2018). We see this very clearly in the sciences. Where once a lone genius could make radical breakthroughs and define an entire field of study, we now require large interdisciplinary and international research teams to make the breakthroughs. Indeed, modern scientific papers frequently seem to require what Daniel Dennett (2017) calls "distributed comprehension", in which no one person understands the entirety of the paper. What all of this means is that achieving greater degrees of rational unity across a human

population has always been essential to our survival and flourishing and has become more essential now. This, again, provides support for the hivemind idea, particularly in its rational unity form. If we want to achieve things through our actions -- for example, maintain the same levels of material wealth and success in the future¹¹ -- and if we want to retain our problem-solving, achievement-oriented society, stronger degrees of rational unity will be required, up to and including degrees that deserve the label 'hivemind'.

One might object to this argument on the grounds that solving problems and achieving goals is of little value unless it is *individual humans* that solve the problems and achieve the goals. This, however, is debatable. Oftentimes the instrumental value of a problem being solved or a goal being achieved massively outweighs the benefit to the individual who solves the problem. For example, it is much more important for our future that a cure for cancer is found, or that our energy or food needs are met, than that some particular individual does so. Similarly, and as stated above, the hivemind ideal may, in any event, be compatible with some lingering individualism that allows for individuals to experience some degree of satisfaction and accrue some benefit from problem-solving and goal achievement. Even if these experiences are not had by an individual human, and instead belong to a collective mind or agent, there is still some value to them: it is still good for certain goals to be achieved and associated experiences to be experienced, irrespective of who exactly is experiencing and achieving the goals. Nevertheless, there may be a nagging suspicion that individualism is somehow implicit in the axiological framework that is at play here—that it is not so easy to dismiss it or brush it off.

The third sub-argument takes a more direct aim at the citadel of individualism and argues that we should favour the hivemind because our commitment to individualism actually undermines truly moral and virtuous behaviour. Impartiality, altruism and selflessness are all usually taken to be hallmarks of moral behaviour. Indeed, a key desideratum in many metaethical theories is the extent to which they warrant impartiality in decision-making (see, for example, discussions of metaethical criteria in Beyleveld 1991, Smith 1994, and Enoch 2011).

¹¹ We pick these as examples of goals that might be deemed “good” and thus add value to life. Some people might disagree with this and think that other goals are good. This doesn’t really affect the point we are making which is simply that some level of goal achievement is part of the good life and many goals require collective action to achieve.

Individualism is an impediment to impartiality. Individualism favours partiality, self-serving bias, and illusions of responsibility and virtue. It is because so many of us are trapped inside an individualistic bubble that we cannot act with true impartiality. Derek Parfit (1984) argues that our belief that we are separate, persistently existing individuals is, to a large extent, a conventional illusion, not a deep metaphysical truth. If we cast off the illusion we can live more open, altruistic lives. As he puts it: “When I believed that my existence was a further fact, I seemed imprisoned in myself...When I changed my view, the walls of my glass tunnel disappeared. I now live in the open air” (1984, 281). In making this observation, Parfit echoes millennia of Buddhist thought, which also emphasises the profound universal empathy that is possible through self-abandonment (Flanagan 2017). In addition to believing that it is a fact that there is no separate, unique, narrative self, members of the Abhidharma Buddhist tradition also think that there are good practical, moral reasons for rejecting an individualistic ethos. As Chadha puts it, they argue that attachment to an individualistic view opens us up to certain “moral defilements” including “greed, conceit, pride, jealousy and so on” (Chadha 2018, 17). We can rid ourselves of these defilements if we loosen the grip of individualism. Furthermore, it is claimed by members of this tradition that our belief in individualism is a source of much of our personal anxiety and suffering. Again, pursuing the hivemind ideal, in both of its forms, would help us to shatter the illusion of independent selfhood and live a more altruistic and enlightened life in the sense that eliminating the distinctions between different agents, and seeing everyone as part of a single unified hivemind, would be to achieve a perfected form of impartial altruism.

It is important that this third sub-argument be fairly interpreted. Staunch anti-individualists and proponents of the No-Self doctrine may push the point too far. While it is undoubtedly true that our commitment to individualism is a source of *some* moral vice and defilement, it is also a source of moral virtue and good (e.g. responsibility, care, duty, courage, charity and so on). A person who is content with the ‘illusion’ of individualism can live a good life. The more modest, and in our view more reasonable, interpretation of the argument is simply to say that individualism is not the only game in town and that there are moral and flourishing modes of existence that are in tension with that paradigm and are best achieved by escaping from it. That is what is made evident by this third sub-argument. That said, there is a significant objection lurking in the background here. One major complaint against this third sub-argument, and as

already noted the second sub-argument, is that they both presume that individualism is indeed an illusion that is easy to cast off. That may not be the case. This is such an important objection, and one that cuts through most of the arguments we are making here, that it is worth addressing separately. We do this in the next section.

The fourth and final sub-argument in favour of the hivemind ideal is that pursuing that ideal is both compatible with, and possibly more conducive to, certain understandings of what it takes to live a meaningful life. This is an argument that has been made by Baptiste Le Bihan in relation to the Buddhist No-Self doctrine (Le Bihan 2018). As he notes, a typical reaction to that doctrine among Western thinkers is to argue that it lends itself to a form of nihilism. If the self does not exist, then what is the point of living? Life is only worthwhile if we, as individuals, live meaningful lives. But, as Le Bihan points out, there is nothing in traditional conceptions of what it takes to live a meaningful life that necessitates individualism. The one exception to this might be certain traditional religious understandings of meaning which insist upon the existence of a separate eternal soul in order for there to be meaning, but this is not a feature of all religious traditions and once you move to more secular understandings of meaning things become much more hospitable to anti-individualist views. The dominant secular understandings of meaning tend to be either subjectivist, objectivist or hybridist in nature (Metz 2013; Danaher 2014 & 2017; Campbell and Nyholm 2015). They claim that in order to live a meaningful life you must either (a) experience some subjective state of well-being (e.g. desire satisfaction or pleasure); or (b) bring about certain objectively valuable states of affairs (e.g. do good, find truth, create beauty); or (c) satisfy some combination of both subjective and objective states (e.g. be subjectively fulfilled by finding out the truth). At first glance, it might seem like each of these theories presupposes or requires the existence of a single self who experiences the relevant subjective states or brings about the relevant objective ends, but this is not necessarily true. There is nothing in subjectivism that insists that subjective pleasure must be felt by a unique, separate and persistent individual, and there is nothing in objectivism that insists that valuable ends must be brought about by a single individual. A collective agent, with high degrees of

rational and phenomenological unity could satisfy these conditions of meaning too.¹² As Le Bihan describes it:

“proper parts of the collections of entities we falsely identify as selves, may collectively instantiate normative properties, and may collectively pursue objective value. The search for meaning is genuine here— but this is not the quest of someone. This is the quest of a plurality that perceives—or thinks of—itself falsely as a singularity.”

(Le Bihan 2018)

Consequently, it is possible to pursue the hivemind ideal without giving up on living a meaningful life—the ideal is compatible with bringing about subjective pleasure and producing objectively valuable ends.

It might even be better than this. There are certain conceptions of meaning in life that are not only compatible with the hivemind ideal but are actually only possible if we pursue it. For example, a common motif in theories of meaning is that a meaningful life is one that contributes to projects or ideals that are beyond that of the individual. The classic version of this is the religious ideal of contributing to God’s universal plan for salvation. One of the more unique secular variations of this ideal comes from Robert Nozick (1981; 1989). Nozick argues that meaning comes from transcending our natural limits (an idea also propounded by Thaddeus Metz (2013) in his theory of meaning), where this requires some connection with larger external things. Nozick argues that there are two ways of forging this connection. You can connect with something that always remains external to you or you can integrate yourself with the external thing, forging some larger organic unity. He thinks the latter strategy is particularly conducive to meaning: it makes us fit as part of a larger pattern. Again, Le Bihan notes that this provides a lot of hope for the proponent of the No Self doctrine because they seek to erase the boundary between the individual human being and the external world. It also provides a lot of hope for the

¹² It is possible that both forms of unity are required on some accounts of meaning but that only one is required on others. For example, on a purely consequentialist theory of meaning -- that meaning depends on producing good consequences for the world -- rational unity may be all that it is required. On subjectivist or hybridist accounts, phenomenological unity may also be necessary.

proponent of the hivemind ideal. After all, what is that ideal if not an attempt to transcend the limits of an individual human body and integrate it within a larger collective unit?

Stepping away from this somewhat mystical doctrine, there is also the common motif that death and the fragility of human life somehow undermines or compromises meaning. Even among avowed naturalists, there is a belief that death makes life less meaningful and at the very least a longer lifespan would be more desirable (Landau 2017). Some people are racing to avoid death by imbibing concoctions of drugs, pursuing scientific research to combat the diseases of ageing and, at the extreme, freezing their corpses in the hopes of later resuscitation (Minerva 2018; O'Connell 2017). There may be something to recommend to each of these strategies, but perhaps there is another way? Perhaps by dissolving the individual and incorporating it into a larger collective agent, we have a more practical route to a longer life? After all, we already know that corporate agents can long outlive their original founders and managers. If we can make ourselves genuinely integrated parts of a unified collective entity we could participate in an extended, possibly even immortal lifespan. This is one thing that the hivemind ideal promises that is not as readily available to individuals.

To briefly summarise, each of these four sub-arguments provides some reason to think that the hivemind society is not just a mere possibility to which we should be open, but is also a possibility that we should desire, because it contains within it a pathway to flourishing.

5. Three Objections and Replies

But there be dragons here. Although we have done our best to argue that the hivemind society is worth taking seriously, we suspect the reader will still have some lingering doubts. We want to close by addressing three that we think are particularly important.

We start with the complaint that it is not that easy to abandon individualism. We alluded to this concern in the previous section when discussing Derek Parfit's views about the shallowness of individual identity and when discussing sub-arguments 2 and 3. It is worth now exploring

those views in a bit more detail. Doing so will enable us to understand why someone might think it is not so easy to abandon individualism, and why we maintain that this is not quite true.

Parfit's views on personal identity are complex, and have been well-ventilated over the years, so there is always a danger that a brief summary does violence to their nuances. Nevertheless, in broad outline, they are as follows. Parfit is a naturalist and reductionist when it comes to individual identity. He does not think that we have a separate, non-material soul that constitutes our individual identity. Rather, he thinks that our individual identities are a consequence of certain functional operations of our brains. He also does not think that there is some deep, naturalistic fact or state of affairs about the brain that determines our individual identity. There is no 'soul-cell' lurking within your neo-cortex, the destruction of which would disrupt or end your sense of personal identity. Instead, your identity is a relatively shallow and contingent fact that results from the way in which your brain instantiates an extended chain of overlapping mental states (beliefs, desires, memories, intentions, etc). In other words, the reason why you think you are the same person right now as you were ten years ago is not because of some deep metaphysical fact linking these two versions of you, but because there is a continuous chain of overlapping mental states linking the present-you to the past-you. That's all there ever is to personal identity. Parfit then argues, through a series of thought experiments, that it is possible to disrupt and play around with those chains of overlapping mental states and so shatter the illusion of continued individual identity.¹³ It is possible to imagine cases (e.g. amnesia-inducing brain injury) where the present-you is not in any meaningful sense the same as the past-you (and ditto for future-you). It is also possible to imagine cases in which the chains of overlapping mental states fission into two streams, both of which can claim the same survival relationship with the pre-fissioning chain, or cases in which the chains are duplicated and cloned. Combined, these thought experiments reveal the relative shallowness of our sense of individuality, and, at least according to Parfit, should encourage us to take a more impartial and altruistic approach to life. There is no really good reason to treat the person who will occupy your physical body in 20 years' time more favourably than the person who occupies the physical body 20 yards down the street.

¹³ To use more Parfitian language, the claim is that the only thing that determines personal survival is the relation between mental states — something that Parfit calls 'Relation R'. This relation comes in degrees and it is possible for you to have relation R with a different physical body.

Parfit's critics are unpersuaded. While many concede that there is no deep metaphysical fact that determines our identities, they argue that there are some robust and resilient practical and biological facts underlying our partiality to ourselves. Two examples of this can be found in papers by Christine Korsgaard (1989) and Tim Christie (2009). Korsgaard mounts a Kantian critique of Parfit's anti-individualism. She argues that our sense of ourselves as unique, enduring agents is not a result of metaphysics but of practical necessity. 'We' each find ourselves situated inside a single body, with different motivations and drives. It is practically necessary for us to resolve these different motivations and drives into a single plan of action, and to do so by occupying a single deliberative standpoint from which we can evaluate and choose different courses in life (Korsgaard 1989, 110-111). Korsgaard accepts that groups of humans can sometimes form group agents, such as nations or tribes (1989, 114) but argues that the existence of such group agents doesn't undermine or eliminate the practical necessity for unity within particular physical bodies. As she sees it:

“when a group of psychological functions occupy the same human body, they have an even more imperative need to become a unified person. This is why the human body must be conceived as a unified agent. As things stand it is the basic kind of agent.”

(Korsgaard, 115).

Christie, whom we encountered earlier, takes a very similar stance, placing particular emphasis on the biological necessity of individuality:

“each human is an individual living organism: humans are physically separate from other things in ways similar to how all organisms are separate from other things. It is empirically false that humans are a ‘super organism’ like, e.g., a hive of bees or a colony of ants”

(Christie 2009, 184-185)

Christie then goes on to argue (2009, 185) that our moral practices are constrained by these natural, biological facts, and hence lead us toward individualism. In combination, both Korsgaard's and Christie's views should give us some pause. Perhaps we are anchored to individualism by practical and biological necessity?

But the biological and practical constraints are not as strong as they initially seem. Four things can be said by way of response. First, as we have already seen, it seems plausible that humans have long been more like bees or ants than we might think since we have already built hives of culture where no individual has a good picture of the final result aimed at by the whole, and the drive for progress is slowly making our endeavors ever more like hives through advances in information communication technology.

Second, *contra* Korsgaard and Christie, there do seem to be at least some human beings who are not so practically or biologically tethered to individualism. Parfit is one. The philosopher Galen Strawson is another. Strawson has written a series of first-personal essays in which he describes his own sense of self (collected in Strawson 2018). He claims that he lacks any strong sense of himself as having a continued identity over time, or of his life fitting within some narrative arc or pattern. Instead, he sees himself as a succession of transient selves, each popping into existence for short periods of time, without any strong connection between them (other than the fact that they share the same biological body). If we assume Strawson's testimony is a sincere and accurate reflection of how he experiences his life, then we have an eloquent and persuasive counterexample to the claim that individualism is a practical necessity. This is to say nothing of the hordes of Buddhists who proclaim their fealty to the No Self doctrine.

Third, the persuasiveness of Korsgaard's and Christie's arguments is entirely contingent on the limitations of the world in which we currently live. Both of them accept this. They argue that the kind of possibility and necessity invoked by their arguments is technological in nature. If technology changed and if we could somehow merge bodies together or erode the practical need for biological unity, then their arguments would no longer hold (Korsgaard 1989, 115; Christie 2009, 188ff). This is an important admission because the argument being put forward in this

paper is concerned with new technological possibilities and how they could facilitate and enable a hivemind society. If we develop technologies for rational and phenomenological unity to a sufficient extent, then it could be the case that what Korsgaard and Christie deem to be practically impossible could become possible, perhaps even normal. A new region in the space of axiological possibilities could be opened up to creatures like us. Christie might argue that even if it did we should resist the temptation to enter that space because of its moral disruption, but the arguments in the preceding section give reasons to doubt this: there is something potentially desirable about migrating into that space.

Fourth, and finally, it is possible to argue that there is no inconsistency between individualism and the pursuit of the hivemind. This is, admittedly, a paradoxical claim. Throughout this paper, individualism has been contrasted with the hivemind, as though they are polar opposites. But if individuals can exist in different sizes—*i.e.* if individuals are not confined to a single human body—then the pursuit of the hivemind society could be seen as the attempt to create a larger, super-individual (*i.e.* group agent or group person). Again, this is not a novel or bizarre claim. The creation of such super-individuals is something that Korsgaard accepts as a genuine possibility when she alludes to the idea of a group agent, and it is something that Carol Rovane has argued for repeatedly over the course her scholarly career (e.g. Rovane 1997, 2004, 2006, & 2014). If the hivemind ideal simply amounts to the desire to create such super-individuals, then many of the values we currently associate with individualism, and many of the things we find attractive about that axiological system, need not be lost in the transition. They simply get transferred to the super-individual.

This brings us to the second objection. This one focuses on the lessons of history. It could be argued that we have tried to create hiveminds in the past and it has not gone well. Consequently, we should not try again. Someone might make the claim that communistic regimes, for example, were driven by something like the hivemind ideal. Individuals were suppressed in the interests of the group or community. The group became the primary unit of moral concern (something similar is probably true for certain imperialistic, nationalistic and aristocratic regimes as well). Historically, this resulted in great misery and suffering. Millions of people were sacrificed in the

interests of the regime, economies were mismanaged, and innovation stagnated. We don't want to go down that road again.

We should certainly proceed with caution and learn from the lessons of history. We should also not coercively apply the ideal to those who are not interested in it or persuaded by its merits: that has been one of the main reasons for historical suffering and immiseration (more on this in a moment). Still, there are reasons to think that both (a) the analogy between the hivemind ideal and historical regimes like communism are not particularly strong and (b) even if they are, the lessons may not apply in the future. On the first of these points, few of the obvious historical analogues involve anything like a pure pursuit of the hivemind ideal. The goal of communists was not to create a single group mind. It was to end capitalism and bring about greater equality and freedom. In many ways it was grounded in a highly individualist ethos. The same goes for nationalistic regimes: they were driven by a desire to create an ethnically pure state that could be glorified and protected from outside invasion. The sacrificial propensity within these regimes oftentimes stemmed from their utopianistic leanings and the associated sense that a few sacrifices on the road to utopia could be tolerated. The pursuit of the hivemind society does not need to be conceived of in such utopianistic ways (it is just one morally desirable possibility among many, not the only one or the best one) nor does it need to be polluted by other goals. Likewise, on the second point, the problems faced by past attempts to achieve hivemind-like status could have resulted from its technological impossibility: people couldn't ignore the practical and biological necessity of individualism. If the technological hurdles to greater collective unity can be cleared, this may no longer be an issue.

All of which brings us to the third and final objection. We just claimed in the previous paragraph that the hivemind should not be imposed coercively upon people against their will. This suggests that we are insisting upon a consent-condition being satisfied before anyone tries to form a hivemind union with another person. It could be argued that a consent condition of this sort is nonsensical in the present context. This is either because consent is a moral constraint that only makes sense in an individualistic world and so if we truly embrace the hivemind ideal

consent is irrelevant. Or it could be because it is impossible for an individual to legitimately consent to form part of a hivemind. It would be like asking someone to consent to be a slave.¹⁴

Neither of these versions of the objection is persuasive. The first version admittedly makes some sense. It could very well be the case that in a world in which everyone has succumbed to the hivemind ideal, consent loses its moral magic and is no longer a relevant consideration. But many of us do not currently live in that world. Many of us live in a world with a strong commitment to individualism. If we are to transition to a different mode of living, then it is plausible to suppose that we should do this by respecting some of the key moral constraints of the world in which we currently live. Consent would be one of them. This means we have to respect consent at least until a significant transition to a hivemind society is made. In other words, at high degrees of rational and phenomenological unity it may well be the case that the individual has no moral standing and consent is irrelevant, but since it is a matter of degree there are, presumably, intermediate forms where individualism and collectivism/hivemindism have some moral parity. In these, intermediate forms the commitment to the consent condition could still be rationally justified, even if only for legacy reasons.

The second version is superficially appealing but it does not withstand closer scrutiny. Transitioning to a hivemind society is not directly analogous to becoming a slave. The problem with consenting to become a slave is that the consent is irreversible and, more importantly, brings about a mode of existence in which you retain your sense of individuality without your freedom. This is part of the tragedy of slavery: that the sense of individuality and autonomy is still present while in the enslaved state. That's not necessarily what happens if you consent to form part of a hivemind. If it is possible to join a form of hivemind society while maintaining sufficient individuality, then there could and should also be freedom for the individuals to *withdraw* (again, assuming there are degrees of hivemind status where individualism either takes precedence over or has parity with hivemindism). On the other hand, if joining a hivemind requires dissolving one's individuality, then a closer analogy might be with consensual

¹⁴ Plausibly the Borg are bad guys because they do *not* seek consent. At least, this would make them bad guys from our current individualistic perspective!

euthanasia or suicide, since they similarly result in the death of the self.¹⁵ In this case, if you think it is possible for someone to rationally and reflectively consent to their own deaths (especially for the sake of something perceived to be of significantly greater value such as an end to suffering), then you should also think it is possible for someone to rationally and reflectively consent to joining a hivemind. That said, it may be wise to insist that people experiment with moderate degrees of the hivemind style of existence before taking any irreversible plunge.

6. Conclusion

In conclusion, far from representing a scary and alien ideal, the hivemind society should be counted among the axiological possibilities for the future of humanity. We are already able to achieve limited forms of hivemind-like status, modest and primitive though they may be, and contemporary technological developments could result in us achieving more impressive forms. When it comes to assessing the potential of these new possible forms of existence we should, at a minimum, adopt a principle of axiological openness and be willing to consider and experiment with them. More than that, though, there are reasons to think that aspects of the hivemind society are desirable. Pursuing it could enable stronger forms of intimacy, problem-solving and goal achievement, moral enlightenment and moral behaviour, while at the same time leaving open most of the traditional paths to finding meaning in life. There are, of course, objections to the hivemind society, but most of these are grounded in concerns that may no longer be relevant in light of technological advances in our capacity to unify with others. Furthermore, no mode of existence is perfect. Everything involves tradeoffs and the prioritisation of some goods over others. The hivemind society is no exception to this. It should be viewed as a legitimate and desirable axiological possibility, but not as the best or only axiological possibility.

¹⁵ They are not perfectly identical of course. In the case of suicide, if you don't believe in an everlasting soul, then the self, and every sub-component or part of the self (including its sentience) ceases to exist. In the hivemind case, some sub-components of the individual will continue to exist, including some aspects of its sentience—they will just be subsumed into the hivemind. If anything, however, this might speak in favour of the permissibility of joining a hivemind since it allows for the persistence of valuable sub-components of the self, and not their everlasting destruction.

References

- Becker, Ernest (1971) *The Birth and Death of Meaning*. 2nd Edition. New York: Free Press.
- Beyleveld, Deryck (1991). *The Dialectical Necessity of Morality*. Chicago, IL: University of Chicago Press.
- Block, Ned (1978). Troubles with Functionalism. *Minnesota Studies in the Philosophy of Science* 9:261-325.
- Bloom, Nicholas, Jones, Charles, Van Reenen, John and Webb, Michael, ‘Are ideas getting harder to find?’ (2018), *NBER Working Paper* 5 March 2018, Version 2.0 – available at <https://web.stanford.edu/~chadj/IdeaPF.pdf>
- Bradford, Gwen (2016). *Achievement*. Oxford: Oxford University Press.
- Campbell, Steven and Nyholm, Sven (2015). Anti-Meaning and Why it Matters. *Journal of the American Philosophical Association*. 1(4): 694-711
- Chadha, Monima (2017). No Self and the Phenomenology of Agency. *Phenomenology and the Cognitive Sciences*. 16: 187-215
- Chadha, Moniam (2018). No Self and the Phenomenology of Ownership. *Australasian Journal of Philosophy* 96(1): 14-27
- Christie, Tim. (2009). Natural Separateness: Why Parfit's Reductionist Account of Persons Fails to Support Consequentialism. *Journal of Moral Philosophy* 6(2): 178-195.
- Churchland, Paul (1981). Eliminative Materialism and the Propositional Attitudes. *Journal of Philosophy* 78(2): 67-90

Clark, Arthur C. (1962). *Profiles of the Future*. Harper & Row.

Collison, Patrick and Nielsen, Michael (2018). Science is Getting Less Bang for its Buck. The Atlantic 16th November 2018.

Danaher, John (2014). Hyperagency and the Good Life - Does Extreme Enhancement Threaten Meaning? *Neuroethics*, 7 (2):227-242.

Danaher, John (2017). Will Life Be Worth Living in a World Without Work? Technological Unemployment and the Meaning of Life. *Science and Engineering Ethics*, 23 :41-64

Dennett, Daniel C. (2017). *From Bacteria to Bach and Back*. W. W. Norton & Company.

Enoch, David (2011). *Taking Morality Seriously: A Defence of Robust Realism*. Oxford: OUP.

Flanagan, Owen (2017). *The Geography of Morals: On the Varieties of Moral Possibility*. Oxford: OUP.

Harari, Yuval Noah, (2016) *Homo Deus*. London: Harvill Secker.

Henrich, Joseph (2015) *The Secret of Our Success*. Princeton, NJ: Princeton University Press 2015.

Huebner, Bryce (2014) *Macro cognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford: OUP, 2014.

Korsgaard, Christine M (1989). Personal Identity and the Unity of Agency: a Kantian Perspective on Parfit. *Philosophy and Public Affairs* 18:101-132

Landau, Iddo (2018). *Finding Meaning in an Imperfect World*. Oxford: OUP.

Le Bihan, Baptiste (2018). The No Self View and the Meaning of Life. *Philosophy East and West* 69(2)

List, Christian and Pettit, Phillip (2011) *Group Agency*. Oxford: OUP, 2011.

Luper, S. (2014). Life's Meaning. In Luper, S (ed) *The Cambridge Companion to Life and Death*. Cambridge, UK: Cambridge University Press.

Maslen, H., SAVULESCU, J. and Hunt, C. (2019) "Praiseworthiness and Motivational Enhancement: 'No Pain, No Praise'?" *Australasian Journal of Philosophy*, pp. 1–15. doi: 10.1080/00048402.2019.1618883.

Metz, T. (2013). *Meaning in Life: an analytic study*. Oxford: Oxford University Press

Metzinger, Thomas (2009) *The Ego Tunnel: The Science of the Mind and the Myth of the Self*. New York: Basic Books.

Minerva, Francesca (2018). *The Ethics of Cryonics*. London: Palgrave MacMillan.

Morris, Ian (2015). *Foragers, Farmers and Fossil Fuels*. Princeton NJ: Princeton University Press.

Nicolelis, Miguel (2011). *Beyond Boundaries: The New Neuroscience of Connecting Brains with Machines and How it Will Change Our Lives*. New York: St Martin's Press.

Nozick, Robert (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Nozick, Robert (1989) *The Examined Life*. New York: Basic Books.

- O'Connell, Mark (2017). *To Be A Machine*. London: Granta.
- Parfit, Derek (1984) *Reasons and Persons*. Oxford: Clarendon Press.
- Paul, LA (2014). *Transformative Experience*. Oxford: OUP.
- Putnam, Robert D. (2001). *Bowling Alone*. Simon & Schuster.
- Rovane, Carole (1997). *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton, NJ: Princeton University Press.
- Rovane, Carole (2004). What is an agent? *Synthese* 140 (1-2):181 - 198
- Rovane, Carole (2006). Why do individuals matter? *Daedalus* 135(4): 49-59
- Rovane, Carole (2014). Group Agency and Individualism. *Erkenntnis* 79 (S9):1663-1684
- Saadia, Manu (2016). *Treconomics*. San Francisco, SF: Pipertext.
- Savulescu, Julian and Persson, Ingmar (2016). Conjoined Twins: Philosophical Problems and Ethical Challenges. *Journal of Medicine and Philosophy* 41: 41-54
- Schneewind, JB (1998) *The Invention of Autonomy*. Cambridge, UK: Cambridge University Press.
- Schwitzgebel, Eric (2015). If materialism is true, the United States is probably conscious. *Philosophical Studies* 172(7):1697-1721.
- Segal, Howard (2012). *Utopias*. Oxford: Wiley-Blackwell.
- Siedentop, Larry (2011). *Inventing the Individual*. London: Penguin.

Smith, Michael (1994). *The Moral Problem*. Oxford: Blackwell.

Sotala, Kaj and Valpola, Harri (2012) 'Coalescing Minds: Brain Uploading-Related Group Mind Scenarios' *International Journal of Machine Consciousness* 4(1): 293-312

Strawson, Galen (2018) *Things That Bother Me: Death, Freedom, the Self, Etc..* New York: New York Review of Books.

Tainter, Joseph (1988). *The Collapse of Complex Societies*. Cambridge, UK: Cambridge University Press.

Warwick, K. (2013). Cyborgs- the neuro-tech version. In Katz, E. (ed) *Implantable bioelectronics — devices, materials and applications*. New York: Wiley.

Warwick, K. (2014). The Cyborg Revolution. *Nanoethics* 8: 263-273.