

INTUITIVE BIASES IN JUDGMENTS ABOUT THOUGHT EXPERIMENTS: THE EXPERIENCE MACHINE REVISITED

Dan Weijers – Victoria University of Wellington

Abstract

This paper is a warning that objections based on thought experiments can be misleading because they can elicit judgments that, unbeknownst to the judge, have been seriously skewed by psychological biases. The fact that most people choose not to plug in to the experience machine in Nozick's famous thought experiment has long been used as a knock-down objection to hedonism. The reluctance to plug in is widely thought to show that real experiences are more important to us than pleasurable experiences. This paper argues that the commonplace choice to remain in reality when offered a life in the experience machine is best explained by status quo bias – the irrational preference for things to remain the same. An alternative thought experiment, preliminary empirical evidence, and a discussion of how psychological biases can affect our judgments are provided to support this argument.

The Experience Machine Objection to Hedonism

The central tenet of traditional Benthamite or Humean hedonism is as follows: only the amount of pleasure and pain we experience (feel from the inside) affect our well-being (Bentham 1789; Hume 1777).¹ This traditional notion of hedonism has been besieged by many hostile arguments, the strength and number of which have led nearly all modern philosophers to believe that it is implausible (Crisp 2006; Feldman 2004; Silverstein 2000; Weijers forthcoming). In particular, an objection based on Nozick's

¹ This kind of hedonism is also called Quantitative Hedonism (Sobel 2002; Weijers forthcoming).

experience machine thought experiment has been cited again and again as evidence that traditional hedonism is “wildly implausible” (Nozick 1974, pp. 42-45; Sobel 2002, p. 244).² This experience machine objection is thought to refute traditional hedonism by casting grave doubt on the intuitive assumption that nothing else really matters to us “other than how our lives feel from the inside” Nozick (1974, p. 43).

In the experience machine thought experiment, Nozick asks us to imagine a machine built by “superduper neuropsychologists” that can provide us with any experience we desire (Nozick 1974, p. 42). This experience machine can provide experiences like publishing an award-winning novel, making a new friend, and many other fantastic experiences that we have never even thought of. Nozick invites us to permanently plug into the experience machine, but expects that most people will forgo a life of great enjoyment in the machine for what they currently have; a presumably less pleasurable but more real life. This choice appears to be based on the “firmly held” negative intuition aroused by considering a life plugged in to the experience machine, an intuition that Sobel thinks strikes “at the heart of hedonism” (Sobel 2002, p. 244).

Nozick proposes that our wanting to *really* experience the limitless reality of living our own real life, as opposed to having a machine live it for us, is probably what prevents us from plugging in to the experience machine (Nozick 1974, pp. 43-45). If this is true, as many people believe it is, then it shows that something like truth or reality matters to us (and not just how pleasurable our lives feel to us from the inside). Based on the inference that truth or reality matter to us, it is concluded that truth (or something like it) must be intrinsically valuable for well-being. The widespread support for this conclusion provides strong *prima facie* evidence that the central tenet of traditional hedonism is false (Weijers 2011).

But is the best explanation for our not wanting to plug into Nozick’s experience machine really that the reality of our experiences matters to us? While the answer to this question might appear to be a resounding ‘yes’, this paper proposes the argument that the choice to forego a life of pleasure

² Examples of authors who have stated or implied that the experience machine thought experiment is a knock-down refutation of hedonism are easily found (e.g. Baggini & Fosl 2007, pp. 74-76; Brink 1989, pp. 223-224; Brülde 2007, pp. 26-29, 33; Darwall 1997, pp. 162, 178; Griffin 1986, pp. 9-10; Kagan 1998, pp. 34-36; Kraut 2007, pp. 124-126; Kymlicka 1990, pp. 13-14; Nozick 1989, pp. 99-117; Sumner 1996, pp. 94-98; van Wyk 1990, p. 109).

in the experience machine is better explained by an irrelevant psychological bias.³

Intuitions and Intuition Pumps

Thought experiments have long been the friend of philosophers, allowing many problems to be addressed without requiring the ‘muddying of knees’ often entailed by field work. More recently, however, warnings have been issued about thought experiments propensity to mislead (e.g. Dennett 1980; Hofstadter & Dennett 1981; Unger 1996; Woodward & Allman 2007). Experimental philosophy has also provided many reasons to doubt the philosophical usefulness of many specific thought experiments over the last ten years (Knobe 2007; Nadelhoffer & Nahmias 2007). This paper continues these emerging traditions using the experience machine thought experiment as a case in point.

Dennett coined the phrase “intuition pump” to describe a thought experiment that (by design or not) elicits a response with a strong intuitive component (Dennett 1980). Typically, judgments about thought experiments will have a deliberative component and an intuitive component (Woodward & Allman 2007). When thought experiments create misleading responses in philosophers, who generally pride themselves on their rational thinking, the intuitive component is more likely to be blamed. Indeed, Bostrom and Ord note that our overall judgments can be “crucially and unavoidably” influenced by our intuitions and the psychological biases they are prone to (Bostrom & Ord 2006, p. 657). Systematic biases also occur in deliberative thinking; however, the combination of two factors characteristic of intuitive judgments makes them more likely to mislead. First, it is very difficult to know when a judgment is misguided and, second, it is even more challenging to ascertain whether the intuitive component of a judgment is tracking relevant, or merely distracting, information. The emerging fields of moral psychology and experimental philosophy have been producing results that are helping us to understand the extent to which intuitions impact our

³ There have been several other attempts to save traditional hedonism from the experience machine objection (Baber 2008; Bronsteen, Buccafusco & Masur 2009; Crisp 2006; De Brigard 2010; Hewitt 2009; Kwall 1999; Kolber 1994; Mendola 2006; Silverstein 2000; Tännsjö 2007). There have also been several attempts to save hedonism from the experience machine objection by changing the traditional account of pleasure (Donner 1991; Feldman 2004; Heathwood 2007; Sumner 1996).

judgments and what features of the thought experiments our intuitive cognition tracks in certain cases. However, this task is made very difficult by the nature of intuitive cognition.

Woodward and Allman provide a neurobiological account of intuitive cognition and contrast it with the other kind of cognition used in judgments about thought experiments – deliberative cognition (Woodward & Allman 2007). They describe an intuition as the visceral sensation that results from a very fast, unconscious and probabilistic processing of many variables in parallel – a definition that is widely accepted in the cognitive sciences (Lieberman 2000; Myers 2004; Woodward & Allman 2007, p. 13). In contrast, deliberative thought is a much slower cognitive process, which consciously uses inductive and deductive reasoning on very limited numbers of variables at a time (Woodward & Allman 2007, p. 13; Bruner 1960). While both modes of cognition have their strengths and weaknesses, and both are susceptible to systematic biases, intuitive judgments have the significant disadvantage of us not being able to know if they have been influenced by various psychological biases (Gilovich, Griffin, & Kahneman 2002). This disadvantage arises because of the process by which intuitions are created.

When novel stimuli are encountered, the brain runs probabilistic inference simulations based on the matches between all of its current stimuli and past experiences (Woodward & Allman 2007). Part of the simulation process involves the reward centre, which produces the message (the visceral sensation or feeling) that we become consciously aware of (Craig 2004; Critchley *et al.* 2004). We are effectively always performing unconscious pattern recognition by monitoring our current environment and comparing it to our archive of experiences. The result of this processing is predictions about what might happen next. These, probabilistically determined predictions are then evaluated using a similar process (comparing them to the value of similar previous actual events) and an overall evaluation is generated. Inferring the value of predicted outcomes in this way can allow for psychological biases to systematically influence our intuitions and, thereby, our overall judgments.

How does this happen with thought experiments? Systematic biases can sneak in during this kind of pattern-recognition process because the features of the current thought experiment (or past experiences to which they are being compared) that have the most weight in the probabilistic

processing might not be the features deemed morally relevant by the readers of the thought experiment. It might be natural to assume that a seasoned philosopher could apply their well-honed rational mind to their intuitions and eliminate any biases before coming to a final judgment. However, even when highly educated (and presumably rational) people are made fully aware of these biases, they still make judgments that bear all of the hallmarks of being adversely affected by them. Consider the self-serving bias: the phenomena that explains why practically everyone (even social psychologists who are well aware of this bias) view themselves as more moral and better at driving than the average person (Myers 2004, p. 95; van Lange, Taris & Vonk 1997). About half of these presumably intelligent people are fundamentally mistaken about some of their own characteristics because of a bias they should have corrected for.

To further complicate the matter, it is difficult to know whether or not a judgment has a strong intuitive component and if that component has been distorted by any biases. Because the causes of our intuitions are processed subconsciously, they are not open to introspection (Lieberman 2000; Woodward & Allman 2007). That is, we might experience a visceral reaction when we come into contact with a new stimulus, such as a thought experiment, but we only really know what the sensation feels like; we know little, if anything at all, about what caused it and why. The best method for establishing the likely causes of an intuition is by reconstruction. By carefully considering all of the possible environmental cues and how they might match a subject's past experiences, one or more of those cues may stand out as obvious candidates for explaining how the intuition was initially constructed. This is the process usually followed (although rarely explicated) by philosophers when they interpret what evidence a particular thought experiment might provide. Using this process, we can reasonably assume that the uneasy feeling we suddenly notice while imagining a thought experiment is probably caused by the thought experiment (unless you have just eaten at that dodgy diner round the corner). What is much less reasonable to assume, however, is which aspect (or aspects) of the thought experiment is causing the intuitive judgment.

Philosophers often assume to have isolated the relevant moral factors in the construction of their thought experiments, but this assumption shows either a misunderstanding of how intuitive judgment works or a lack of awareness that intuitive judgments affect our overall judgments. The

pared-down, and often unrealistic, thought experiments used by philosophers usually stipulate only very few aspects of the situation under assessment. Indeed, the decision to create thought experiments in this way is often justified by claiming that less background information prevents irrelevant factors from being considered. This is, unfortunately, not necessarily the case. Recall that intuitive cognition operates by comparing the new stimuli with existing experiences to try to predict what might happen next. Minimalistic thought experiments are matched to the closest real experiences. Those real experiences will not be minimalistic, however; they will contain many features absent from the thought experiment, including emotionally salient (and therefore powerful) ones. Subconscious simulations are then run based on the real experiences that the thought experiment was most closely matched with. Since this process takes into account features that were relevant to the real experiences, but not included in thought experiment, the resulting intuitions will be partly based on irrelevant information.

Making matters even worse, however, is the fact that most philosophical thought experiments stipulate features that are so unrealistic that we have not experienced anything like them – indeed we are likely to have experienced the very opposite of them. When these clashes occur, our intuitions are likely to be based on information that is not just irrelevant, but contrary to the point of the experiment itself.

It is these features of intuitive cognition that enable structural biases to affect our judgments about thought experiments in ways that they do not affect our deliberative cognition.

The Trip to Reality

When asked to consider Nozick's experience machine thought experiment, the vast majority of people make the judgment that they would prefer their current life over a life hooked up to the machine (Nozick 1974, pp.42-45). This judgment is likely to have a strong intuitive component, since it is not always obvious to people (when first exposed to the thought experiment) why they would not prefer a life of pleasure in the experience

machine.⁴ The majority of reasons provided by philosophers for preferring their own life are, in line with Nozick's justification – that a life in the experience machine is not real and that experiences based on reality are more valuable for our well-being than those created by a machine (Nozick 1974, pp. 43-45). How can we assess whether the falseness of the experiences in the experience machine is really what makes us judge that our current lives are more valuable? By considering similar thought experiments that only change allegedly irrelevant factors, and then observing our judgments in those cases, we can see which factors best explain the resulting judgments (Kolber 1994). The following thought experiment, the Trip to Reality, holds constant the realness of experiences inside and outside of the machine, while changing a few other purportedly irrelevant factors.

Imagine that you leave your family for a weekend to attend a conference on the experience machine thought experiment. While you are there, someone informs you that you are actually in an experience machine. She offers you a red and a blue pill. She explains that taking the blue pill will take you to reality and taking the red pill has the dual effects of bringing you back from reality and completely wiping any memories of having being in reality or being offered funny pills. Being a curious philosopher you swallow the blue pill. It turns out that reality is quite different to the world you have been experiencing inside the machine. Some of the most noticeable differences are that your experiences are a more mundane and less enjoyable than before. You also discover that nearly all of your friends and family are either in experience machines or do not exist in reality! Your father is there, so you spend time with him. But, a few conversations reveal that he is not really the person you know as 'Dad'. It is time to make the choice. Will you take the red pill so that you can go back to your previous life (with no idea that it is not in fact real)? Or will you throw the red pill away and try to make the best life you can in the more real, but less enjoyable, surrounds of reality?

In the Trip to Reality thought experiment, the thought of getting into an experience machine does not elicit the same intuition that exactly the same act does in Nozick's experience machine thought experiment. In my

⁴ Since deliberative judgments are open to introspection, and intuitive judgments are not, judgments that appear to be formed *because of* a reason are more likely to have a large deliberative component and judgments that appear to be formed without any immediately obvious reason are more likely to have a large intuitive component.

experience of presenting the two scenarios, dramatically more people choose a life in an experience machine when considering the Trip to Reality thought experiment than when considering the experience machine thought experiment. Initial empirical results from experimental philosophy endorse this claim.

Philippe de Brigard conducted several small studies using various twists on Nozick's experience machine thought experiment (de Brigard 2010). Unsurprisingly, participants responding to a vignette similar to, but much more exaggerated than, the Trip to Reality nearly unanimously opted to spend the rest of their lives in an experience machine.⁵ In this vignette, their real life was much worse than their current life (they were a prisoner in a maximum security prison in real life). But, very surprisingly, participants responding to a vignette, in which their real life was being a multimillionaire artist living in Monaco, were divided equally between choosing reality and the life they had in an experience machine!⁶ This result implies that people's opinions are divided on whether an unreal version of their current life is better for their well-being than a life that sounds great *and* would be real, but is unfamiliar.

De Brigard conducted a follow up study to further investigate what might be affecting the participants' choices (de Brigard 2010, p. 49). De Brigard compared two neutral vignettes and found that adding the statement "your life outside is not at all like the life you have experienced so far" to one of them made the corresponding participants much less likely to choose reality.⁷ The sample sizes were relatively small in de Brigard's experiments, but the results are significant enough to consider his justification for the divergence in responses between the choice of getting into an experience machine and the choice of getting *back* into one; namely, status quo bias.

⁵ Only 3 (13%) of the 24 participants answering this vignette opted for reality (de Brigard 2010, p. 47).

⁶ 12 (50%) of the 24 participants answering this vignette opted for reality (de Brigard 2010, pp. 47-48).

⁷ In the first neutral vignette, 13 (54%) of the 24 participants opted for reality (de Brigard 2010, p. 48). In the second neutral vignette only 33 (41%) of the 80 participants opted for reality (de Brigard 2010, p. 49). So, adding the statement "your life outside is not at all like the life you have experienced so far", made 13% less respondents choose reality. De Brigard did not use this statistical method to evaluate his results, but applying a two-tailed Fischer's Exact test to his data reveals that we can be about 91% confident that the difference in responses to these two vignettes is not the product of chance.

Status Quo Bias

A group of overlapping psychological biases, best referred to as status quo bias, provide the best causal explanation for both the judgment that a life in Nozick's experience machine is worse than real life and the judgment that reality is worse than a life in an experience machine in the Trip to Reality thought experiment. These overlapping psychological biases are importantly linked by our valuing prospective gains only about half as much as we value avoiding equivalent prospective losses of things we already have or know – the status quo (Gilbert 2006; Gilovich, Griffin & Kahneman 2002; Kahneman & Tversky 2000; Kahneman, Knetsch & Thaler 1991; Samuelson & Zeckhauser 1988; Tversky & Kahneman 1991). Simply put, status quo bias is an irrational preference to keep things the way they are. Such preferences are considered irrational because they assign value to certain things over and above any utility value they might have (in the broadest possible sense).⁸

The endowment effect, an aspect of status quo bias, has been used to explain why only about 10% of undergraduate students, who were rewarded with either a mug or a chocolate bar for filling out a survey, took up the cost-free opportunity to swap their reward for the other type⁹ (Knetsch 1989; Knetsch & Sinden 1984). An irrational preference for the status quo has also been posited as the explanation for Hartman, Doane and Woo's field study of power consumers (Hartman, Doane & Woo 1991). The consumers were sorted into two groups, one for consumers with *more* reliable and expensive power services and one for consumers with *less* reliable and expensive power services. When provided with six reliability-to-cost mixes, with one option indicated as their status quo, the vast majority in both groups expressed a preference for the mix indicated as the status quo for their group (60% and 58%) while only a tiny fraction wanted the reliability-to-cost mix that the other group had (both 6%) (Hartman, Doane & Woo 1991, p. 149). Possible explanations, other than status quo bias, do exist for both of these examples,

⁸ Note that even items that are monetarily worthless can have immense utility value. Take, for instance, your child's first pair of baby shoes. Their resale value is approximately nil, but every time you see them, you recall pleasant memories that bring you joy. An irrational preference for these shoes would be to prefer them to something else that would give you more joy (all other things being equal).

⁹ Cost-free is perhaps not totally accurate here, since participants had to raise a sheet of coloured paper that was on their desk to indicate that they would like to switch rewards. However, this cost would be so small that it would only affect the behaviour of the laziest students... most of which would probably still have been in bed during the experiment.

but the literature in support of status quo bias is extensive and, while the mechanisms underpinning status quo and other related biases might not yet be fully understood, psychologists are generally convinced that the effects of status quo bias on our judgments are significant and widespread (Gilovich, Griffin & Kahneman 2002; Kahneman & Tversky 2000).

In the experience machine thought experiment the choice to get into the machine involves giving up something very important and (in nearly all cases) valuable that you have and are familiar with (your current life) for something that is supposedly more valuable but fairly unknown to you (a life in the machine). It is clearly risky to consider swapping your current life for another one if you are unsure of what the new life will be like. I propose that both this caution regarding the unknown and an irrational over-valuing of the familiar is affecting intuitive judgments regarding the experience machine thought experiment. Evidence for this can be found by a comparison with the Trip to Reality thought experiment, in which being in the machine is framed as our current (familiar) life. This comparison amounts to what Bostrom and Ord call the reversal test, which specifically assesses whether status quo bias is an important causal factor in the resulting judgment (Bostrom & Ord 2006). The reversal test turns the scenario around to frame the supposedly operant variable as the status quo rather than with a change in circumstances. In the Trip to Reality thought experiment, being in an experience machine is described as our current real life, which means the status quo for us is being hooked up to a machine. Comparing both scenarios, we discover that the majority of judgments are in favour of maintaining the status quo *regardless of how real the future experiences would be*.

We must pause at this point to consider if there are any differences between the original experience machine and the Trip to Reality thought experiments, other than the framing of what is the status quo. It turns out that there are other differences that seem to affect our judgments in these cases, but that they are all related to the status quo. The most notable difference is that the risks involved in each case seem markedly different. In the original experience machine thought experiment, our intuitive cognition would have deemed a machine life as risky, despite the stipulation in the thought experiment that the machine works perfectly. This intuition of risk likely arises from all of our previous experience with computerized machines crashing at least once, if not regularly, and often not providing the quality of

performance that they promise.¹⁰ In the Trip to Reality thought experiment, the risk of machine failure and machine underperformance are unlikely to affect our intuitive judgment about a life in a machine because that scenario would be matched to our non-crashing real-life experiences during the intuitive processing of the thought experiment.

It might be argued that these considerations imply that a preference for the status quo is not at all irrational; why would it be irrational to prefer to avoid potential risk? Of course, taking risk into consideration would not be irrational if you were offered to connect to an experience machine in real life. It would be rational to demand the highest possible level of evidence that there was no risk – something that is rarely possible in real life. In Nozick's original thought experiment, however, the risk of machine underperformance or failure was ruled out by stipulation. Therefore, considering these risks during judgment formation about the thought experiment is irrational; it gives weight to irrelevant factors.

Another factor that seems likely to affect our judgments in the original experience machine thought experiment is Nozick's stipulation that we are to ignore family considerations (Nozick 1974, p. 43). Nozick explains that we do not have to worry about this because they could all plug into a machine as well (Nozick 1974, p. 43). It might be argued that the choice not to plug into a machine is therefore based on the rational preference not to force such a risky decision on your whole family, or if risk could somehow be eliminated, the rational preference not to force such an important decision on your whole family. This is a defect in Nozick's formulation of the thought experiment – it brings an important real-life factor into consideration that should be omitted when arguing about well-being, or what the good life *for* the person living it is. A person's machine life could be experientially identical, in respect to experiences with family, to a real life and need not include the experience of forcing them to do anything. If Nozick had simply stipulated that considerations of obligations to others should be ignored, then considering them would become irrational; it would give weight to irrelevant factors.¹¹ It is less obvious that the preference to maintain our current family relationships as they are is related to status quo bias, but it does reflect the most important aspects of an irrational preference for the way things are. The

¹⁰ Note that this is my attempt at reconstructing the intuition and, for the reasons discussed above, is far from a precise process.

¹¹ In view of the discussion above, however, it should be clear that our intuitive cognition would probably not be able to comply with the stipulation to ignore our family.

preference would be irrational in this case because it is for the real family connections over the different (but experientially the same or better) family connections that would result in a machine life. Recall that when in the machine we would not know that we were no longer experiencing reality.

In a similar vein, it might be argued that the choice to remain in the machine in the Trip to Reality thought experiment is the result of a rational, as opposed to irrational, preference for the close relationships and other goods that we have and know in our machine life.¹² I agree, since this argument undermines the experience machine objection to hedonism. The original experience machine thought experiment was devised to show that reality (and not just how our experiences feel on the inside) should matter to us and, since traditional hedonism does not intrinsically value reality, it must be false. The argument that the Trip to Reality thought experiment reveals that it is in fact rational to prefer the life we experience over a *real* life, is great news for traditional hedonists.

All things considered, there are several potential differences between the experience machine and Trip to Reality thought experiments, and the reasons why they might elicit the judgments that they do. All of the important differences are related to the status quo, however; they are all related to what we have, what we know, and how we value those things. Therefore, the best explanation for why most people prefer their real life to a life plugged into Nozick's experience machine is not that real experiences are important for our well-being. Rather, the best explanation is that people's judgments are heavily influenced by an irrational preference for what they already have and know – by status quo bias. This, of course, leaves room for parts of the explanation to remain missing, but not important parts.

Conclusion & Implications

What does the foregoing discussion mean for Nozick's experience machine thought experiment? It must no longer be given as evidence that real experiences are more important to us than pleasurable ones. Taking this change in orthodoxy seriously means that the experience machine objection cannot rationally be used to show that pleasure and pain are not the only factors that contribute to well-being and, thereby, that hedonism is false.

¹² I thank an anonymous reviewer for drawing my attention to this point.

Hedonism, however, is not magically saved from the many criticisms it faces just because the experience machine objection to it is misguided. Other evidence is available to support the proposition that pleasure and pain are not the only ultimate determinants of well-being. For example, Shelly Kagan's deceived businessman thought experiment addresses the question of false pleasures and seems to provide evidence that pleasures based on deception are not as valuable as pleasures based on truth (Kagan 1998). Furthermore, the deceived businessman thought experiment appears to do this without incurring any distortion from status quo bias. The more that is learnt about how our judgments are formed, however, the more certain kinds of thought experiments might be shown to be irrelevant or misleading, perhaps including Kagan's deceived businessman.

Nevertheless, that the experience machine thought experiment has misled the widespread and firmly held judgments of so many philosophers stands testament to the fact that we should all be less confident about using thought experiments *alone* as evidence for our arguments. This is especially the case when those thought experiments are unrealistic because our intuitive cognition is much more likely to process irrelevant factors in such cases. The status quo bias is of course just one of many psychological biases that can have major effects on any of our judgments with a significant intuitive component. So, heed the warning; many thought experiments in philosophy should be revisited by applying what we are currently learning about intuitions to ensure that more of our firmly held judgments are not, unbeknownst to us, marred by the effects of psychological biases.¹³

Bibliography

Baber, H., 2008: 'The Experience Machine Deconstructed' in *Philosophy in the Contemporary World*, 15(1), pp. 132-137.

Baggini, J. & Fosl, P.S., 2007: *The Ethics Toolkit: A Compendium of Ethical Concepts and Methods*. Wiley-Blackwell.

Bentham, J., 1789: *An Introduction to the Principles of Morals and Legislation*. Adamant Media Corporation, 2005.

¹³ Many thanks to Nick Agar, Bob Goodin, and an anonymous reviewer for insightful comments on earlier drafts.

- Bostrom, N. & Ord, T., 2006: 'The Reversal Test: Eliminating Status Quo Bias in Applied Ethics' in *Ethics*, 116, pp. 656-679.
- Brink, D., 1989: *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Bronsteen, J., Buccafusco, C. & Masur, J., 2010: 'Welfare as Happiness' in *Georgetown Law Journal*, 98, pp. 1583-1641.
- Brülde, B., 2007: 'Happiness Theories of the Good Life' in *Journal of Happiness Studies*, 8, pp. 15-49.
- Bruner, J., 1960: *The Process of Education*. Cambridge, MA: Harvard University Press.
- Craig, A.D., 2004: 'Human Feelings: Why are Some More Aware than Others?' in *Trends in Cognitive Sciences*, 8, pp. 239-241.
- Crisp, R., 2006: 'Hedonism Reconsidered' in *Philosophy and Phenomenological Research*, LXXIII(3), pp. 619-645.
- Critchley, H.D., Wiens, S., Rotshtein, P., Öhman, A. & Dolan, R.J., 2004: 'Neural Systems Supporting Interoceptive Awareness' in *Nature Reviews Neuroscience*, 7, pp. 189-195.
- Darwall, S., 1997: 'Self Interest and Self Concern' in *Social Philosophy and Policy*, 14, pp. 158-178.
- De Brigard, F., 2010: 'If You Like it, Does it Matter if it's Real?' in *Philosophical Psychology*, 23(1), pp. 43-57.
- Dennett, D., 1980: 'The Milk of Human Intentionality' in *Behavioral and Brain Sciences*, 3(3), pp. 428-430.
- Feldman, F., 2004: *Pleasure and the Good Life*. Oxford: Clarendon Press.

- Gilbert, D., 2006: *Stumbling on Happiness*. New York: Alfred A. Knopf.
- Gilovich, T., Griffin, D.W. & Kahneman, D., 2002: *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press.
- Griffin, J., 1986: *Well-Being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon Press.
- Hartman, R.S., Doane, M.J. & Woo, C.-H., 1991: 'Consumer Rationality and the Status Quo' in *The Quarterly Journal of Economics*, 106(1), pp. 141-162.
- Hewitt, S., 2009: 'What Do Our Intuitions about the Experience Machine Really Tell Us about Hedonism?' in *Philosophical Studies*, 151(3), pp. 331-349.
- Hofstadter, D. & Dennett, D., 1981: *The Mind's I: Fantasies and Reflections on Self and Soul*. New York: Basic Books.
- Hume, D., 1777: *An Enquiry Concerning the Principles of Morals*. T.L. Beauchamp (ed.), Clarendon Press, 1998.
- Kahneman, D., Knetsch, J.L. & Thaler, R.H., 1991: 'Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias' in *The Journal of Economic Perspectives*, 5(1), pp. 193-206.
- Kahneman, D. & Tversky, A., 2000: *Choices, Values, and Frames*. Cambridge: Cambridge University Press.
- Kagan, S., 1998: *Normative Ethics*. Boulder: Westview Press.
- Kawall, J., 1999: 'The Experience Machine and Mental State Theories of Well-being' in *The Journal of Value Inquiry*, 33, pp. 381-387.

- Knetsch, J.L., 1989: 'The Endowment Effect and Evidence of Nonreversible Indifference Curves' in *American Economic Review*, 79, pp. 1277-1284.
- Knetsch, J.L. & Sinden, J.A., 1984: 'Willingness to Pay and Compensation Demand: Experimental Evidence of an Unexpected Disparity in Measures of Value' in *Quarterly Journal of Economics*, 99, pp. 507-521.
- Knobe, J., 2007: 'Experimental Philosophy' in *Philosophy Compass*, 2(1), pp. 81-92.
- Kolber, A., 1994: 'Mental Statism and the Experience Machine' in *Bard Journal of Social Sciences*, 3(3-4), pp. 10-17.
- Kraut, R., 2007: *What is Good and Why: The Ethics of Well-Being*. Harvard University Press.
- Kymlicka, W., 1990: *Contemporary Political Philosophy: An Introduction*. Oxford: Oxford University Press.
- Lieberman, M., 2000: 'Intuition: A Social Cognitive Neuroscience Approach' in *Psychological Bulletin*, 126, pp. 109-137.
- Mendola, J., 2006: 'Intuitive Hedonism' in *Philosophical Studies*, 128(2), pp. 441-477.
- Myers, D.G., 2004: *Intuition: Its Powers and Perils*. Yale University Press.
- Nadelhoffer, T. & Nahmias, E., 2007: 'The Past and Future of Experimental Philosophy' in *Philosophical Explorations*, 10(2), pp. 123-149.
- Nozick, R., 1974: *Anarchy, State, and Utopia*. Oxford: Blackwell, 1991.
- Nozick, R., 1989: '10. Happiness' in his *The Examined Life*. New York: Simon and Schuster, pp. 99-117.

- Samuelson, W. & Zeckhauser, R., 1988: 'Status Quo Bias in Decision Making' in *Journal of Risk and Uncertainty*, 1, pp. 7-59.
- Silverstein, M., 2000: 'In Defense of Happiness: A Response to the Experience Machine' in *Social Theory and Practice*, 26(2), pp. 279-300.
- Sobel, D., 2002: 'Varieties of Hedonism' in *Journal of Social Philosophy*, 33(2), pp. 240-256.
- Sumner, L., 1996: *Welfare, Happiness, & Ethics*. Oxford: Clarendon press.
- Tännsjö, T., 2007: 'Narrow Hedonism' in *Journal of Happiness Studies*, 8, pp.79-98.
- Tversky, A. & Kahneman, D., 1991: 'Loss Aversion in Riskless Choice: A Reference-Dependent Model' in *Quarterly Journal of Economics*, 106(4), pp. 1039-1061.
- Unger, P., 1996: *Living High and Letting Die*. New York: Oxford University Press.
- van Lange, P.A.M., Tarris, T.W. & Vonk, R., 1997: 'Dilemmas of Academic Practice: Perceptions of Superiority among Social Psychologists' in *European Journal of Social Psychology*, 27, pp. 675-685.
- van Wyk, R.N., 1990: *Introduction to Ethics*. New York: St. Martin's Press.
- Weijers, D., forthcoming: 'Hedonism' in *The Internet Encyclopedia of Philosophy*.
- Weijers, D., 2011: 'The Experience Machine Objection to Hedonism' in *Just the Arguments*. Michael Bruce & Steven Barbone (eds.), Wiley-Blackwell.

Woodward, J. & Allman, J., 2007: 'Moral Intuition: Its Neural Substrates and Normative Significance' in *Journal of Physiology-Paris*, 101(4-6), pp. 179-202.

Dan Weijers

Philosophy Department, Victoria University of Wellington,

PO Box 600, Wellington 6140, New Zealand

danweijers@gmail.com