



THE UNIVERSITY OF  
**WAIKATO**  
*Te Whare Wānanga o Waikato*

Research Commons

<https://researchcommons.waikato.ac.nz/>

## Research Commons at the University of Waikato

### Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

**The Virtual Society:  
An examination of the concerns regarding a life in  
fully immersive virtual reality**

A thesis  
submitted in fulfilment  
of the requirements for the degree  
of  
**Master of Social Sciences in Philosophy**  
at  
**The University of Waikato**  
by  
**MAXWELL RODLEY**



THE UNIVERSITY OF  
**WAIKATO**  
*Te Whare Wānanga o Waikato*

2025

## **Abstract**

This thesis argues that in the future, fully-immersive virtual societies could contribute positively to the prudential wellbeing of their citizens, even if they are aware that their reality is simulated. It builds upon the work of David Chalmers in *Reality+* (2022), by further detailing how a virtual society might address the current trepidations many people have about living a life in virtual reality. Additionally, the thesis formulates a coherent model of a prudentially valuable virtual society, a model that can be critiqued and refined over time.

The first section describes several components of life in non-virtual reality that are considered by skeptics to be lacking or non-existent in virtual reality. For each example, it is argued that a virtual society can find a means to incorporate them. The second section discusses the position that utopian societies are "too good" to be prudentially valuable, and provides counterarguments against some of the reasons behind that perspective, in the context of a virtual utopia. Lastly, the third section offers five 'axioms of virtual abuse prevention'; design objectives that may help protect a prudentially valuable virtual society from corruption, abuses of power, and the suffering that may result.

## **Preface**

For a long while, I have thought about the ways societies could be arranged to help people lead better lives. Whenever I hear of voting systems that better represent the preferences of citizens, or learn of scientists working towards cures for aging, or discover videos about human-friendly urban planning, it sparks a great excitement in me. Not only for the fact that so many possibilities exist, but knowing that diligent and persuasive people are working to bring such ideas from the drawing board into the new normal. My motivation behind this thesis is to help spread more positive possibilities, both my own and those of others, and hopefully encourage more people to think about the kind of future they want to see.

## **Acknowledgments**

My deepest thanks to my supervisors Dan Weijers and Nick Munn, for your patience, advice, and your valuable insights.

And thanks to my family and friends for your love and support, both in the real and digital realms.

## **Table of Contents**

Abstract.....	ii
Preface .....	iii
Acknowledgments.....	iii
Introduction .....	1
Importance of this discourse .....	1
Nozick’s Experience Machine.....	2
Gaps in Reality+.....	3
Overview .....	4
Section 1 – Concerns of Wellbeing .....	5
If one is aware they are in a simulation, there is no deception.....	5
“Simulated experiences are not authentic” .....	7
‘Placehood’: permanence and impact in virtual worlds .....	11
“Simulations cut us off from reality” .....	12
The Gut Feeling Problem .....	15
Summary of Section 1 .....	16
Section 2 – “Too-Goodism” .....	17
An Unearned Reward.....	18
Boredom and Monotony .....	18
Overabundance.....	22
Haidt’s Anxious Generation .....	25
Summary – What is the Challenge? .....	28
Section 3 – Abuses of Power .....	28
AVAP 1 – Disincentivization .....	29
AVAP 2 - Immutable Safeguards .....	31
AVAP 3 - Numerous Stakeholders and Governmental Structure .....	33
AVAP 4 - Transparency and Privacy .....	36
AVAP 5 - Emergency Measures.....	38
Summary of Section 3 .....	39
Conclusion.....	39
Summary .....	39
Potential Implications .....	39
Areas of Further Inquiry.....	40
References .....	41

## **Introduction**

In his book *Reality+* (2022), Chalmers makes the case that simulated realities are real, and therefore simulated experiences are real too. This axiom is then used to help justify further claims throughout the text. One particular claim is that if future simulations become “...indistinguishable from the nonvirtual world” (Chalmers 2022, pg. 8), it would be possible for someone to live a good life inside of one. Additionally, Chalmers distinguishes these simulations from ones such as Nozick’s experience machine in order to alleviate some concerns that arise from that thought experiment. Lastly, Chalmers uses an objective-list approach to well-being in order to identify sources of value that he argues are still present in perfect simulations.

My thesis likewise aims to prove the persistence of a series of prudential contributors to well-being in the context of Chalmers’ simulations (what he calls the ‘reality machine’), particularly those of concern to philosophers such as Nozick. Some of these contributors will be those that Chalmers has touched upon in *Reality+*, and these will be dissected in greater detail. Two important facets I will refer to when considering these contributors are the awareness that one is in a simulation, and the assumption that simulated realities are not ‘real’ metaphysically, as Chalmers (2022) argues. These facets are taken into account in an effort to bolster the opposing arguments. In addition, I will argue against the position that simulations abundant with fulfilling activities are ultimately detrimental to living a good life, a position I will call ‘too-goodism’. Lastly, I will consider abuses of power in a virtual society and how those concerns may be alleviated.

## **Importance of this discourse**

To some, it may seem impractical to consider versions of virtual reality (VR) that are yet to be realised. Lassiter and Kagan (2022, pg. 7) argue that Chalmers’ reality machine hypotheticals are “largely speculative “futurism” which tends to veer into hand-waving and waxing philosophical”, and that philosophers should “focus on “imperfect” and “partially immersive” virtual realities [so we may] apply these insights and suggestions today”.

I agree that current VR is rife for philosophical discussion, but the truth of that matter does not make discussions of perfect VR any less important. Considering a possible outcome for a technology has its value in the here and now. If planners in the 1890s had considered what a fully car-centric transport network might be like and opted to avoid that future, our streets might be much safer for pedestrians today. A less speculative example is mutually assured destruction, a concept highlighted by the likes of Alfred Nobel in the 1870s (Tägil 1998) and crystalized following the development of nuclear weaponry. If there is a way to begin construction of a utopia, we ought to consider if it even is one before we begin. Otherwise, we may waste our efforts or overlook a better alternative. We are not yet at a point to say whether or not we can build a perfect virtual reality, but I think it best for us to have an answer on whether we should before that point arrives.

Of course, we should consider VR as it exists today when discussing its future. As of now, we do not know if the virtual realities of today and the automobiles of the 1890s are similar in terms of future adoption. That being said, there is considerable interest in the space. According to a Grand View Research report (2023, n.p.), “The global virtual reality (VR) market size [...] is expected to grow at a compound annual growth rate (CAGR) of 27.5% from 2023 to 2030”. This comes in the wake of new virtual and augmented reality headsets continuing to be released, such as the Meta Quest 3 and Apple’s Vision Pro, both announced less than two years ago (Chin 2023; Pierce 2023), and the Meta Quest 3S releasing in the latter half of last year (Hunter 2025).

Compare the investment into virtual worlds and the investment into bases and colonies in outer space. Multiple countries and agencies are seeing renewed interest in space exploration and settlement, in particular NASA's Artemis program that plans to develop permanent structures on and above the moon's surface (Davenport 2023). SpaceX CEO Elon Musk also remains determined to establish a Martian colony (Wattles 2023). If it is important to consider the value and social structure of larger-scale extra-terrestrial colonies, given that plans are underway to construct small ones, then it is also important to consider a potential grander future for our current virtual colonies.

## Nozick's Experience Machine

It is worth noting the differences between the kind of simulation Chalmers and myself are exploring (what Chalmers calls the 'reality machine') and the one Robert Nozick (1974) explores in *Anarchy, State, and Utopia*. These differences not only alleviate some of the concerns Nozick and others have with a simulated life, but also create new concerns. Nozick's experience machine has the following properties:

1. While you are connected to the machine you are living a simulated life, separate and distinct from a 'real' life.
2. Your experiences and actions in the machine are predetermined and scripted.
3. You are unaware that you are in an experience machine.
4. No one else is present with you in the simulation, only simulated humans.
5. You cannot choose to exit the experience machine once you are in (instead determining the amount of time you spend within before entering).

The reality machine shares only that first property with the experience machine. The remaining properties are no longer applicable. In regards to the second point, Chalmers (2022, pg. 306) notes that "VR is not preprogrammed. Typically, it's open-ended. A user in the reality machine exercises choice, and what happens there depends on the choices the user makes. [...] so users can indeed be genuinely courageous or genuinely kind in the reality machine".

When we consider the reality machine as an advanced version of the VR we have today, the remaining aspects of the experience machine are also of no concern to users of the reality machine. In modern VR, you remain aware that your experiences are simulated, you can interact with other human beings, and may return to reality whenever you wish. Chalmers' reality machine functions in the same vein, albeit with more advanced hardware. Chalmers argues that these properties help alleviate Nozick's concerns where the experience machine only exemplified them, particularly concerns of agency and to a certain extent accomplishment.

However, the reality machine raises different concerns. The first is the 'utopian' aspect. Users are aware that they live in a simulation, have memories of their life prior, and can move in and out of the simulation at will. As such, living out a simulation similar to an ordinary life (even a rather lucky one) would probably not be sufficient for them. Chalmers (2022, pg. 348-349) notes this, stating that:

"Virtual worlds may remove scarcity of many material goods. Space is not at a premium in VR. Everyone can have a personal idyllic virtual island if they choose to [...] The result may be virtual abundance [...] Under virtual abundance, important material goods in virtual worlds are instantly reduplicable and available to all. "

Considering our limited resources in the physical world, this level of abundance would be very tempting. So even if reality machines did not provide it initially, it's likely that alternatives would,

jumping ahead of the competition as a result. And this abundance extends beyond material goods too. In the reality machine, where all senses are engaged with the virtual world, practically any experience would be possible, from flight to a fully immersive Dungeons & Dragons roleplaying campaign to blowing up the moon. The concern here is 'too-goodism'; the idea that this degree of choice and possible experiences may not be good for us in the sense that it is 'too good', and that limitations are a necessary evil.

Another concern is the distinct social aspect of reality machines. In an experience machine, users live in wholly separate simulations from one another. Meanwhile the reality machine "allows your friends and family to share a reality with you" (Chalmers 2022). One can imagine that these virtual worlds would grow to facilitate larger and larger communities. A modern example of this would be VRChat, a series of virtual spaces both communal and private that hosts thousands to tens of thousands of players at any given time (Steam Charts n.d.).

This scenario raises the important question of how communities like this ought to function. Unlike real world communities, these virtual communities are immediately and existentially dependent upon the maintenance and security of real-world software and hardware. Even if a town experiences a blackout in the real world, residents can still interact, exchange, and live in that community. If multiple citizens of a virtual community lose access to the internet or electricity, that community ceases to be accessible for those citizens for as long as the power is out. If the servers and their backups fail, then all of its citizens are cut off.

Therefore, when considering virtual communities, the means of sustaining those vital real-world systems, and who or what owns those means, is as crucial as food and water. This is on top of various other distinctions between virtual and real communities, such as a lack of resource scarcity, the greatly reduced need for labour as a result, and the open question of how one might raise a family in virtual reality, if at all. If one is to have a good life in a virtual society, that society needs to be good in turn. Figuring out if a healthy and stable virtual society is possible, therefore, is vital.

## **Gaps in Reality+**

There may be an unspoken agreement between Chalmers and Nozick in Reality+. Chalmers (2022, pg. 305) believes that one of Nozick's worries about simulations is their illusory nature, and that because of this, one "merely [has] the experience of writing books and making friends" without actually doing anything. Chalmers responds to this concern by arguing that simulations are not illusions, meaning that real actions are possible in the reality machine (Chalmers 2022). This raises the issue that, if he is wrong about simulations being real, then this response is insufficient in alleviating Nozick's worry.

I am uncertain whether or not Chalmers agrees with Nozick that someone could not actually do things in an illusory simulation. As noted earlier, he believes that "a user in the reality machine exercises choice" (Chalmers 2022, pg. 306). However, the ambiguity and the reliance on simulations as non-illusory is a potential weakness in Chalmers' broader argument that one can live a good life in the reality machine.

Beyond this, there are a few important prudential and practical factors to wellbeing that are not explored in depth and in some cases unexplored. Another of Nozick's worries is a lack of "contact with a deeper reality", which Chalmers interprets as one that is not built by humans. Chalmers responds by pointing out that many of us already live within cities; themselves highly artificial environments. This is a good point, but I doubt it would be enough to alleviate Nozick's worry here.

First of all, even within our artificial habitats, one can still perceive aspects of the natural world like the sky, wind, oceans, rivers, and life. More importantly, buildings and settlements are a part of our physical reality. Each city has an impact on other cities and the wider world, and vice versa. If it is prudentially important for us to remain cognizant participants in this physical system of interactions, then there is a clear difference between living in a city and living in a simulation, regardless of artificiality.

As for gaps in the more practical and instrumental factors to wellbeing, privacy is a notable one. In his chapter on building a virtual society, Chalmers (2022, pg. 347) asks “Are the corporations that run virtual worlds manipulating our behavior? Are they invading our privacy?”. He then speculates that in the future, inhabitants of virtual worlds will push for and create more fair and equitable virtual worlds, ones that presumably do not invade people’s privacy.

There is very little here on the specific issue of privacy in virtual worlds, and the solution proposed is not very useful. While it is possible that independent virtual worlds could become freer and more utopian, we are also seeing worlds emerge that are far from that ideal. One such case is the virtual world of Decentraland. Its user-led government is described by Folding Ideas (2023, n.p.) as having “no actual authority” and is “speedrunning bureaucratic paralysis”. According to his report, “[a] committee intended to have minimal discretion is completely unrestrained with zero accountability”.

These are just a couple of examples to illustrate the gaps present in Reality+. For each prudential and practical contributor to wellbeing that I discuss in this thesis, I will highlight the manner in which Chalmers addresses it.

## Overview

I will be dividing my arguments across three different sections. The first section will address several of the concerns often raised by skeptics of permanent life in virtual reality; specifically concerns related to the wellbeing of its inhabitants. Some of these concerns will be ones touched upon in Reality+, which will be expanded upon further. The concerns I will address are the notion of virtual reality as ‘deceptive’, a lack of authentic accomplishments in virtual reality, a lack of permanence and personal impact, and a disconnect from reality. I will argue that these concerns, while applicable to certain forms of virtual reality, are not intrinsic to virtual reality. I will conclude this section by addressing the underlying sense of unease many feel about living in virtual reality, the possible causes, and how it may be diminishing over time.

Section Two will expand on the virtual society by envisioning a virtual utopia, and the potential concern that such a society would be ‘too good’, in the sense that a society with too much excess and too little risk would be bad for one’s well-being. There are four aspects to this that I will explore; the sense that one has not ‘earned’ a utopian life, a ‘too good’ society leading to boredom and monotony, the risks of overabundance leading to desensitization and decision paralysis, and the negative effect such a society could have on children and their development. Similarly to Section One, much of my argumentation will involve alleviating these concerns, and also highlighting the ways a virtual utopia can improve upon current non-virtual reality in these areas.

Section Three will be about ensuring that a virtual society can maintain its beneficial contributions to well-being, and not fall victim to tyranny or abuse. I outline five ‘Axioms of Virtual Abuse Prevention’ (or AVAPs) that a virtual society can adopt to aid in this goal. The first involves identifying the motivations for both seeking and abusing power, and how a virtual society can mitigate them, both through active measures and through the way the society is organized. The second AVAP involves

the measured application of immutable safeguards to ensure political flexibility while also bolstering democratic resilience against tyranny and corruption. The third centres around how the political structure of the virtual society can play a major part, and argues that a combination of direct democracy and sortition would further support democratic resilience and participation. The fourth AVAP would explore how a virtual society may handle privacy and government transparency, and the final AVAP would cover emergency measures in the event of virtual abuse bypassing the four previous AVAPs; the measures divided between prevention of democratic breakdown, and ‘last resort’ measures focused on citizen evacuation and safety.

Over the course of these sections, we will gather an understanding of what a prudentially positive virtual society might look like, as certain design features help solve multiple concerns at once, and are able to operate in tandem with other components.

## **Section 1 – Concerns of Wellbeing**

When presented with the concept of life in a virtual society, there are many reasons one might have to reject the idea. One of those reasons is the belief that a virtual society cannot replace certain aspects of our current lives; aspects that are considered valuable to our wellbeing (Nozick 1974, pg. 43; Chalmers, 2022, pg. 305-306). These aspects may include genuine connections with others, not being deceived by reality, and the satisfaction of doing ‘meaningful’ activities. Any of these might also be accompanied by a more instinctual trepidation; a feeling of unease about plugging in permanently. The aim of this section is to highlight the potential to address each of these concerns, regardless of their validity, when constructing a virtual society. This will serve to indicate that a virtual society is not damaging to prudential wellbeing simply due to being virtual.

### **If one is aware they are in a simulation, there is no deception**

Chalmers (2023, pg. 306) believes that simulated worlds and objects are not illusions because they are real worlds and objects, and that because of this fact, there is no deception in his reality machine. I would add that even if we accept the notion that simulated objects are illusory; in the sense that they provide similar sensory information as a real object would without actually having the properties of that object, illusions are not necessarily deceptive.

This is not to say that illusions cannot be deceptive. Within our non-virtual reality, we believe the causes of our qualia are tied to specific external processes, such as light reflecting off a tree and into our retinas, giving us the experience of a tree. If it were somehow revealed to us that our non-virtual reality was an illusory reality such as a computer program, and that there never was a tree, we could rightly claim that we were deceived, since the sources of our qualia were different than what we were led to believe; computer code instead of physical objects.

However, if I go into a simulation with the awareness that I am doing so, and I remain aware of the simulation throughout my time within it, then I always know the true source of my qualia. I know that my experiences in the simulation are caused by a computer program, and when I exit the simulation, I understand that the cause of my qualia has shifted to physical processes. Unlike in the previous example, I was never led to believe that the cause of my experiences was different to what it was. The same is true when witnessing optical illusions like Kaniza’s triangle, or playing a computer game like Grand Theft Auto. Despite perceiving a white triangle, or a major traffic violation, I am not being tricked into thinking that my perception stems from a real white triangle or a real traffic

violation. Therefore, just as one can experience an optical illusion or play a video game without being deceived, one can live in virtual reality without being deceived.

Additionally, virtual realities and video games are far from the only illusions in our lives. Our perceptions are biased by our biology and past experiences. If a baby, an adult, and a snake attended a university lecture, the adult might hear a literary dissection of *Moby Dick* while the baby would hear a semirhythmic string of mouth noises. The snake would hear even less, but would get a great look at the lecturer's body heat. It is hard to say that any of the observers in this scenario are witnessing the whole truth of the matter. Rather, each of them experiences individual aspects of the whole, while also being blinded to the extent of that whole truth. The baby knows that a human is making noises, but is blind to their meaning. The adult is focused on the meaning of those noises, the words they form, without parsing them as the series of sounds that they are. And both humans are ignorant of how the lecturer's body heat shifts and grows, unlike the snake.

The illusion here, and in real life, is that what one is perceiving is the whole picture. This is an illusion that everyone falls for most of the time out of practical necessity, except when one is reminded of it. At best, we are forced to suspend our disbelief of the idea that our perception is wholly accurate, in order to live our lives. This is also the case in VR. In order to engage with the virtual, we must suspend our disbelief of the virtual being real. In fact, there is more certainty regarding the nature of reality itself in virtual reality compared to non-virtual reality, since you know that everything you experience is an illusion, or at least parsed through an illusory medium. This is not the case in non-reality, as everything you perceive may be a truth, a half-truth, or an illusion.

Of course, there are still ways of being deceived in a simulation. You could chat with an artificial intelligence (AI) and believe that you are talking to a fellow human being. But this problem is not unique to virtual reality. One news story noted a Google employee believing that an artificial intelligence he was talking to was "self-aware" (Guardian staff and agency 2022). The question therefore becomes a matter of deceptive quality; whether or not deceptions are easier to achieve or convince more people in the reality machine than in our non-virtual reality. This is not an easy question to answer, as we do not know how the reality machine is structured in terms of hardware, software, or the finer details of the experience.

However, there is also no reason to assume that VR would not also become a means of calling truth to power, or that malicious spreading of misinformation would run rampant without being restrained or corrected. Such corrective measures were already being undertaken across the internet. Prior to the current Trump administration, services like X (formerly known as Twitter), Google, YouTube, and TikTok were taking steps to limit the spread of misinformation, providing corrections when it does appear, and being held to account in this regard by governmental bodies like the United States Senate and the European Commission (Asselin 2023).

Of course, it is highly debatable whether or not we can trust tech companies to combat misinformation effectively without oversight. At the same time, rebuilding public trust in honest experts and accountable institutions is not impossible, nor is fostering an appreciation for proper research and fact-checking in our education systems. There is no reason why we could not take these same steps in virtual reality. Any deception we may experience in a virtual world is not the inevitable result of that world being virtual, nor the result of inevitable mass misinformation.

## **“Simulated experiences are not authentic”**

In *Anarchy, State, and Utopia* (1974), Nozick notes a distinction between experiencing a thing and doing a thing. He claims that one reason someone ought not plug into the experience machine is because:

“We want to do certain things, and not just have the experience of doing them. In the case of certain experiences, it is only because first we want to do the actions that we want the experiences of doing them or thinking we’ve done them” (Nozick 1974, pg. 43).

A distinction between experiencing and doing is reasonable to me. If I launched a VR program that gave me the sensory information of Ringo Starr playing on the Ed Sullivan Show, I would only be experiencing the sensation of drumming without actually doing any drumming. Even in a more interactive example, such as playing a modern shooter game, I would not actually be firing a gun. I would only be pressing a button to create a simulacrum of a gun firing. However, I do not believe that this distinction justifies someone not plugging into a reality machine.

Firstly, just because someone is not actually doing what they are experiencing does not mean that that person is doing nothing. In the case of the shooter game, I am using my reflexes in an attempt to complete an objective; beating the game. When experiencing Ringo’s senses, I am gaining an understanding of the situation and entertaining myself, just as one might do when watching a documentary. However, if this second form is all that the experience machine is; a collection of experiences with no need or means to do anything but experience them, then I agree with Nozick. As he puts it, “Plugging into the machine is a kind of suicide” (Nozick 1974, pg. 43). Just as we need to do more than simply watch television or TikTok all day in order to live a good life, we need to do more than passively experience a simulation. We need to do things that have authenticity; in the sense that the actions we take should reflect our own efforts and desires in each moment, rather than being predetermined. Fortunately, there are already real examples of authentic activities occurring in virtual worlds.

An anthropological study of the online virtual reality game VRChat found a “wide variation of interactions between persons” (Montemorano 2020, pg. 35). Notably among those interactions were very open and deep discussions about real-world events, as Montemorano describes:

“The red-haired avatar and the robot avatar [...] converse intensely about their lives outside of this reality, discussing happiness, identity, and taking life for granted. They reassure each other, thoughtfully offering advice and examples as the conversation shifts fluidly from one topic to another” (Montemorano 2020, pg. 36).

This is far from an isolated case. A YouTube search for ‘VRChat interview’ yields dozens of videos of individuals discussing their traumas, mental disorders, breakups, and more, all conducted and recorded in the virtual space. It is not unreasonable to assume that hundreds or even thousands of these earnest conversations have taken place during VRChat’s history. In addition, the virtual nature of the space could very well be fostering more earnest discourse from people who otherwise would keep to themselves. In the YouTube video ‘Identity, Gender, and VRChat’, an anonymous user describes how “I’ve never felt comfortable with who I am [but] in VRChat, I just became more comfortable with myself over time” (Straszfilms 2021, n.p.). In that same video, the author; who is well-acquainted with the VRChat community, notes that:

“Our behaviour, our identity, is influenced by the people that we are around and the groups that we belong to. [...] It makes sense then [that] an entirely different group [...] separate

from reality [is] going to have a profound impact on people's identity and how comfortable they are about sharing bits of themselves that perhaps they never have before.”

Based on these experiences, it seems that authentic conversations are a reality in virtual spaces. However, these conversations are still dependent upon other real beings to deliver their authenticity, much like how a phone requires multiple people to enable a long-distance conversation. In scenarios wherein interacting with real beings becomes difficult, one might be forced to interact with simulated humans instead. If we assume that these simulated humans are similar to philosophical zombies in that they are not conscious and considerate but merely give the illusion that they are (Chalmers 1996), then it seems dubious to claim that such interactions are authentic. However, this problem is not unique to virtual reality.

Services like Character AI enable users to interact with chatbots that emulate the personalities of various people, both real and fictional (Gunnell 2022). These emulations can already produce some compelling conversations. When I asked an AI version of Socrates if he believed that interactions with chatbots could be authentic, he noted that “Both parties can learn from each other and gain new perspectives on various topics, which is the essence of authenticity”. If technologies like this become even more advanced and ubiquitous, more and more people may turn to them as sources of advice and companionship, regardless of how virtual reality technology develops. With that being said, virtual reality may contribute to their popularity. In a virtual space, artificial characters could be given a body and the ability to interact with users kinaesthetically. Instead of just talking to an artificial Superman, you could also share a meal with him and fly around New York together.

Even though we are assuming such interactions are not authentic, perhaps we can use them to guide users towards more authentic experiences. Since interacting with other users in virtual space is authentic, artificial characters may be designed to provide opportunities to talk to real people for those with problems communicating. Our AI Superman could invite you to play tennis with genuine people, and as you grow more comfortable around those people, the need to interact with Superman decreases. Alternatively, the creators of virtual societies may choose to limit the scope of their artificial characters, making them more akin to non-player characters in video games than multifaceted beings.

Of course, there is more to an authentic life than authentic socialisation. There are also things we wish to authentically do in our lives. This is why receiving praise for an accomplishment, like building a model ship, is usually more satisfying when you have genuinely accomplished that task instead of, for example, paying someone else to build the model ship for you. These authentic tasks include ambitious goals like summiting mountains or writing novels, as well as the more day-to-day activities that make us feel productive and valuable, like helping others with problems or preparing meals. Such activities are prudentially valuable for our well-being, and for many there is an intuitive difference between cooking food and climbing mountains in non-virtual reality compared to cooking and climbing in a simulation.

Some, like Nozick (1974, pg. 43), may view that simulated mountain climb as simply the experience of climbing decoupled from the ‘doing’ of a genuine climb. Others may see the achievement of climbing a virtual mountain as lesser than climbing a real one, either due to the relative reputations of the two mountains, or a perceived easiness or limitation of the virtual ascent. I will offer a contrary perspective, using mountain climbing as a stand-in for the wide range of achievements possible in a reality machine.

First of all, what makes climbing a mountain valuable in the first place? Those who argue the prudential value will cite accomplishment, that one has accomplished something in reality, while

climbing a virtual mountain is not an accomplishment, or at least not as substantial of an accomplishment. I fully agree that accomplishment, tremendous effort, and satisfaction have value to wellbeing, and I would argue that all of those things could be found in summiting virtual peaks to an equal if not greater degree.

Consider a more gradual transition between a genuine climb and a simulated one. Forcehimes and Semrau (2016) offer a few transitional steps; including a climb with your brain inside a prosthetic body almost identical to your natural one (which they call Prosthetic Unity), and climbing with a prosthetic body that you control remotely while your brain is in a vat (which they call Remote). Their argument is that if your experiences, challenges, and decisions are the same in each scenario, and since you are ultimately still pulling a physical body to the top of a mountain in each scenario, then there should not be any difference in your well-being, or in the amount of accomplishment you experience (Forcehimes & Semrau 2016).

This is further exemplified in the 'Uneventful Ascent', 'Blackout' and 'Flicker' scenarios. These cases match Prosthetic Unity, only the prosthetic also sends information about the climb to an experience machine, which your brain connects to if there's a loss of connection between your brain and the prosthetic (Forcehimes & Semrau 2016, pg. 463). During these connection gaps, your brain experiences a "counterfeit but identical experience" through the experience machine while the prosthetic continues to climb based on your inputs. In 'Blackout', a substantial portion of the climb takes place in such a gap. In 'Flicker', multiple intermittent gaps occur throughout, while in 'Uneventful Ascent', no gaps occur at all.

Forcehimes and Semrau (2016, pg. 464) find it hard to accept that Blackout is worse than Uneventful Ascent, or that compared to Blackout, one "enjoys a higher level of well-being in Flicker", simply because of the differing levels of contact with reality in each case. I am very much in agreement with this line of reasoning. From the 'Blackout' and 'Flicker' scenarios, it becomes less of a stretch to apply the same logic to climbing a mountain in virtual reality; the only significant difference is the location the climb is taking place in.

To exemplify this argument, instead of copying a real mountain, let us build a new one for our virtual world. Mount Ultimus has the difficulty of K2 and twice the height of Everest, with a unique topographic profile to navigate. And while those two other mountains have had hundreds of summiteers by now, you are the first to try and summit this digital behemoth and must plan for everything yourself.

Since Ultimus is in the reality machine, all of your senses are incorporated into the struggle, including the aching of your muscles and the bite of the cold wind. At every moment of the climb, you make choices just as freely as you would in reality, and a mistake might cost you the climb. And while death may not be on the table should you fall, one could easily imagine other consequences. Maybe Ultimus only gives you one attempt and if you fall, you lose access to it for a time, possibly indefinitely, or you could lose the use of a limb for subsequent climbs, or maybe the simple prospect of starting back at the bottom is sufficient danger.

What Ultimus shows is that challenges in virtual worlds can be unique, demanding, engaging, and even more challenging than those in non-virtual reality. Struggling to the top of Ultimus would be a monumental achievement, just as it is to summit K2 or Everest. Even if the mountain is illusory, the struggle and willpower required is not. And unlike Everest, there are no established camps, litter, summit queues or dead bodies. Nothing to detract from the experience of summiting a gruelling mountain beyond what you take with you. You could even do so with a group of friends and bond through the task.

One might argue that summiting Ultimus does not have the same value as summiting K2 because of the difference in reputation. K2 has a long and deadly history for climbers, and its reputation stems from that, as does the recognition that comes from submitting it. Meanwhile, Ultimus may just be one virtual mountain of millions. Putting aside the fact that wholly private accomplishments can still be rewarding, digital obstacles are not immune from garnering considerable reputations. A salient example of this can be found in the video game speedrunning community.

In speedrunning, your goal is generally to complete a specific game as quickly as possible. As faster and faster records are set, the methods required to set a competitive time increasingly rely on luck, new discoveries, and precision throughout each run (Platts 2020). One game where precision is particularly important is *Getting Over It*, where you are tasked with using precise movements and tricks to climb a mountain. Less than 10% of players have managed to complete it (Steam Community n.d.), garnering the game a considerable reputation for its difficulty (Wood 2017; Ditzler 2018).

Ever since people first began speedrunning *Getting Over It* in late 2017, not a single player out of thousands had managed to complete the game in less than a minute, even by the start of 2023 (Speedrun.com n.d.). Eventually, a few months later, this challenging task was finally accomplished by a player named Blastbolt (Blastbolt the Bolt 2023). Over a million people have watched his run on YouTube, with many commenters describing it as one of the most significant achievements in speedrunning history.

What this example demonstrates is that challenges in virtual environments can not only be difficult, but can garner considerable reputations and draw massive amounts of attention to those who meet them. Therefore, Mount Ultimus could definitely garner a reputation, as would anyone able to summit it. Its climbers may not be climbing a real mountain, but they are still “doing” significant things. This “doing” is also true of the more day-to-day activities one might undertake in the reality machine, like preparing a meal. Every step of the cooking process could still be required, and in a simulation as advanced as the reality machine, the quality of your work on each step would create the exact same dish as it would in non-virtual reality; an exceptional shepherd’s pie, a modest shepherd’s pie, a burnt pie, or anything in between.

As a result, making a good shepherd’s pie ‘legitimately’ in the reality machine would be just as difficult as making one in non-virtual reality. ‘Legitimately’ in this case describes the undertaking of tasks in the reality machine in the same manner as you would in non-virtual reality, instead of simply snapping your fingers to make a shepherd’s pie appear. And, with the difficulty and time required to legitimately make a shepherd’s pie in the reality machine, there would also be the satisfaction of seeing it through, and the joy offered by the process itself and the learning of new skills.

Even if one considers simulated worlds to be unreal, there are still activities one can genuinely do in a simulation. Many writers today use computers not only to write books but to distribute them to be read on other computers. Even if those books never sell physical copies, most people would accept that their writers have written books. If a writer followed this same process while living in a virtual world, the only thing that has changed is the medium. Unlike in the experience machine, where only the experience of writing is granted to the user (Nozick 1974), the writer in the reality machine would still be choosing which words to write, just as climbers on Ultimus choose which rocks to grab. Additionally, those writers would still be responsible for planning out their characters, their plots, and how their books will be structured, just as modern writers do. If someone can genuinely write and sell books on a computer, then someone could genuinely write and sell books in a reality machine.

Regardless of what we wish to do in life, plugging into a reality machine does not compel us to simply experience life, without making meaningful choices or putting in effort. In the reality machine, each of us is free to choose what we do, who we do it with, and how much effort we put into it. The machine enables one to authentically socialise, accomplish, and live.

## **‘Placehood’: permanence and impact in virtual worlds**

One notable aspect of non-virtual reality is that actions cannot be undone, and those actions impact both the present and future. This serves to make decisions important, as the consequences of someone’s decisions affect the rest of their lives and the wider world in varying degrees. Part of the appeal of older cities like London and Rome is that the structures that make them up were all built at different times by different people, giving each place a visible history and culture. As Chalmers (2022, pg. 315) puts it; “There’s value in visiting places where history-shaping events happened. There’s value in taking part in time-honoured traditions”, yet our modern virtual worlds are often more “transient”. In the reality machine, you could build Rome in a day, smash it to pieces, and bring it back unscathed, leaving no trace of the destruction.

Therefore, if these simulated worlds are so malleable, people may feel like they cannot leave a lasting impact, or that their decisions do not particularly matter. At the same time, a significant appeal of reality machines is the freedom to experience almost anything, which necessitates a significant degree of malleability. Fortunately, there are ways to balance malleability and societal impact, depending on how a virtual society prioritises public and private spaces.

Perhaps each individual has their own mini-universe, and these universes are linked together by a common space, resembling a Victorian town, the Enterprise (from Star Trek), or anything else. This ‘communal’ space must be passed through in order to visit other mini-universes, and might be less alterable compared to the mini-universes, or perhaps changing it requires collective agreement. The virtual society could have dozens of distinct common spaces, or there might be a vast common space encompassing dozens of different communities across incredible distances. I will explore this last option, as it seems to best address the issue of impact for the following reasons;

1. The presence of a community space that has a (simulated) physicality to it, and the necessity of entering that space in order to travel and socialise, better fosters awareness and appreciation of that space as opposed to making it an optional destination. As such, incorporating all of the community spaces in the society into one vast area would better foster awareness and appreciation of the society as a whole, making it feel more like a living breathing nation. At the same time, the variable distances between different communities can help maintain community identity and lead to interesting social dynamics between those communities.
2. I noted that alterations to such spaces would be more limited compared to each person’s more malleable mini-universe. In combination with Reason 1, this could make each change a person or group of people makes more meaningful, since it is less easily undone and anyone in the nation could find and enjoy it for years to come.

This solution addresses not only the concern of personal impact, but also of virtual locations not having the same history as non-virtual ones. Chalmers (2022, pg. 315-316) notes that “no virtual world we’ve yet created has anything like the long history of our nonvirtual world, and history is something that many of us greatly value”, yet he also acknowledges that “virtual worlds will presumably have notable histories of their own in the long run”. By creating a more limited but still

somewhat malleable overworld, we can provide a space for that notable history to take place and be visible to all.

One might argue that even if people need to pass through this overworld to visit each other's universes, most will see it only as a transitional space. They might claim that there is less incentive to interact with your friends in a limited overworld compared to interacting with them in the more limitless personal spaces, where any activity is accessible. This is a valid concern. Even with our current technology, many people are opting to spend the vast majority of their time within their homes rather than in the wider community. Japan even recognizes such behaviour as a health condition; hikikomori (Kato et. al. 2019). This then begs the question of how such lifestyles are addressed in society, which methods are successful, and whether or not those methods can be translated into the reality machine.

Answering this question is beyond my expertise and the scope of this thesis, but there are some potential psychological benefits to the overworld that the personal universe might not necessarily provide. These unique benefits make the overworld prudentially valuable, and incentivises interaction with it. Firstly, with the overworld being less malleable, it could help provide a sense of stability. It would be the primary place to forge new friendships, and for those who grow up in the reality machine, it may be the source of their first friendships. Unlike visiting a friend's universe, where they have ultimate power over the experience, the overworld would provide ways to socialise on more equal footing. This practice may well extend to larger scale events, such as sports or concerts. Attending such events, alongside limitless exploration of the overworld's various communities and environments, could provide the satisfaction of being socially and physically active in a rich and complex world, made as such by the individuals that inhabit it.

Virtual communities can be 'bustling'; filled with people and opportunities for socialisation across different groups. VRChat, World of Warcraft, Roblox, and Zepeto, are examples of this, with particular spaces becoming hubs for interaction between friends and strangers alike (Montemorano 2020; Williams et. al. 2006; Oh et. al. 2023). These spaces may even aid in overcoming social anxieties and introversion, becoming more enticing to socially anxious individuals than non-virtual spaces as a result (Ayobi, Deighan & O'Kane 2023).

Not every virtual community is like this of course. Decentraland is a salient example; the only location where players can be found consistently is filled with idle users and robots waiting to click on meteors in the hopes of acquiring non-fungible tokens in order to sell them off later for a believed profit (Folding Ideas 2023). However, in some virtual spaces, it is possible to develop bustling locations of general socialisation, despite there being more secluded locations available. Given what is possible in the reality machine, replicating such places seems feasible, giving citizens a reason to participate in the wider community and become part of its history. So, the concern of permanence and impact, while valuable to consider when designing a virtual society, is not a fundamental issue with the concept.

## **“Simulations cut us off from reality”**

Nozick's states that plugging into a simulation would “limit us to a man-made reality [with] no actual contact with any deeper reality” (Nozick 1974, pg. 43). One way of interpreting this statement is that our lives would have no meaningful impact, as the last section explored. Another way of interpreting this statement could point towards Nozick being concerned about artificiality and information. 'Artificiality' in the sense that any simulation is entirely man-made with no genuinely natural components. 'Information' in the sense that simulations may keep us sheltered and ignorant to the

truth; either in a conspiratorial sense or in a deeper sense of understanding reality that only experiencing it can provide. For example; understanding the mechanisms of colour is considered by some to be different from genuinely seeing and experiencing colour (Jackson 1986).

While one could dismiss the concern of artificiality as a fallacious appeal to nature; where what is natural is considered good by default (Curtis 2020), it may also be the case that those living in very artificial environments like cities are more at risk of both mental health issues and reduced appreciation towards the natural world (Gruebner et al. 2017; Turner, Nakamura & Dinetti 2004). However, this possible link does not implicate artificiality necessarily; only the specific artificial environment of the city. And while the world of a reality machine is technically more artificial than a real-world city; since genuine natural elements like the sky or the sea would be absent, the natural element of other human beings would still be present. Humans are not artificial, and as established earlier, natural human beings can still interact and bond in a simulated environment.

Additionally, virtual worlds may help curb those aforementioned mental health risks. Since virtual worlds are infinitely malleable, an infinite variety of environments are possible to live in, including those that may feel more natural to us than real-world cities. Virtual communities could live in treehouses in a dense rainforest, on the side of a mountain rich in cherry blossoms, or in the early American midwest. These small communities would be surrounded by interactive environments that feel just as real as the real thing. If experiencing more natural environments in smaller communities is beneficial for one's mental health, then living in virtual communities like these may be better than living in a real-world city where nature is necessarily curbed and a sense of community may be harder to find.

What is less clear, however, is whether or not an appreciation for such communities would translate into a greater appreciation for the non-virtual natural world. If we become enraptured by our virtual environments, it is not unreasonable to predict that we could come to neglect the non-virtual world in the process. That being said, people living in reality machines would still have reasons to care about the non-virtual world. First of all, such simulations are dependent upon hardware and other infrastructure in the real world, meaning that residents of the simulation would have a vested interest in protecting that infrastructure, along with ensuring that their physical bodies are kept healthy. These bodies and systems would require nutrients, water, oxygen, electricity, and other resources.

Unsustainable resource extraction eventually leads to resource scarcity, and instability increases the risk of system failure. Assuming that individuals living in the simulation want to maintain the simulation for as long as possible, it would make sense for them to favour means of extracting resources that are sustainable and do not lead to instability. In addition, there will undoubtedly be people who choose to live in the non-virtual world regardless of how well-off those living in the simulation might be. Communities that avoid certain technologies exist today; the Amish and Mennonites being well-known examples. A simulated society will need to take into account their relationship with such groups, especially if those living in simulated worlds have friends or family who do not. Otherwise, the simulated society risks conflict and instability both external and internal.

If good relations are to be maintained, then the simulated society needs to consider the real-world locations that their allies rely on and take steps to avoid damaging those locations. This in turn would improve the reputation of that society and increase positive interactions. Individuals might feel more confident migrating into the simulated society knowing that they could talk to, visit, or even migrate back to their old communities if they wished. If the simulated society has excellent relations with its neighbours, agreements regarding mutual defence and maintenance could also be enacted, further increasing stability. Taking these factors into consideration, I believe that those

living in a simulated society would still have a vested interest and concern for the non-virtual world, even if we assume that their exposure to virtual nature would not itself foster an appreciation for non-virtual environments.

There is still the matter of information to consider. If a simulated society is steeped in deception, its government could take advantage of citizen's lack of experience of the non-virtual world to push a particular narrative. For example, the government might claim that the non-virtual world has been rendered uninhabitable in order to justify creating systems that prevent citizens from leaving. In our current societies, such claims could easily be debunked by the public through simple observation of the world around us. A simulated society curates everything that each individual experiences, so a deceptive virtual society could simply prevent any such observations from taking place, replacing them with evidence that justifies their claims. A citizen might ask the simulation to display a live video feed of the complex that houses her body in the non-virtual world, and the simulation could alter the footage to feature long-dead flora, skeletons, and a toxic yellow sky.

In order to alleviate issues of deception, a virtual society would require a means of deterring deceptive practices. Imagine a conniving leader wishes to alter the code simulation-wide such that unflattering images of him could never be accessed. Maybe any changes to the simulation-wide code would be automatically communicated to every citizen in simple language, with each citizen being allowed to accept or decline as they wish. Perhaps any decisions made by leadership must also be accepted by an ethics committee appointed by and from the citizenry. In addition, those living outside the simulation may catch wind of unethical behaviour and take action against the leadership. I will discuss the specifics of regulation in a later section, but these examples serve to illustrate that virtual societies can arise with both internal and external factors that deter deception.

However, one might argue that even a simulated society without deception would be ignorant to the experience of living in non-virtual reality. While we have discussed the idea of citizens leaving the virtual society to visit or even immigrate to the non-virtual world, there might still be a difference in perspective between those who travel to the non-virtual world and those who are born and raised there. The same is true today of people who move to another country and have to adjust to a new way of life, in comparison to those born in that country who have a deeper understanding of life there. In the case of non-virtual and virtual living, there would be some relevant distinctions. The non-virtual world may not be as abundant with experiences, nor would it be as malleable as virtual life. Experiences such as pain, exhaustion, and hunger would not be optional. It is reasonable to fear that growing up without these constraints in a virtual society may be so different from normal human development as to be detrimental.

If this is the case, then perhaps the children of a virtual society need to live in 'maturity boxes'; simulated worlds that are accessible by their primary caregivers and function as reasonably safe mirrors of reality. This way they can develop skills and overcome challenges just as one would in non-virtual life, growing and maturing as their ancestors did. In order to prepare them for the virtual utopia that awaits them, maturity boxes might mirror a point in history where discussions of potential virtual societies are becoming more and more common. We might be in a kind of maturity box ourselves in that regard. Once a certain age or level of maturity has been reached, those inside are able to slowly transition out of their maturity box and live in the virtual society proper.

One of the big questions the maturity boxes pose is whether or not the children inside them will be made aware of the virtual society that awaits them in the future. On one hand, children being aware of an inevitable virtual paradise may throw a significant spanner into their natural development. They might be less inclined to work hard or think of their actions as meaningful. On the other hand, some might see the mass concealment of the virtual society as too great of a lie to be ethical.

However, it is common and even expected for parents to conceal aspects of the world from their children, such as swearing, violence, and pornography. Maybe maturity boxes and parents would adopt a more gradual approach, slowly revealing more and more about the virtual society as each child matures.

Regardless, there are many reasons why a virtual society would still be aware of, understand, appreciate, and have access to the outside world. So long as the mechanisms to prevent obfuscation persist, this understanding will persist too.

## **The Gut Feeling Problem**

Despite the reassurances that I have laid out thus far, some may still feel uneasy about living in a reality machine, even if they consider my arguments to be reasonable. Their reasons may be tied to concerns of governance or of a simulation being 'too perfect', both of which we will explore in later sections, but beyond that there is a general wariness of simulated life in the public sphere. I have discussed this topic with a number of individuals from multiple age groups and professions, and while some of them were keen to live in a reality machine, there were also several people who remained hesitant even when their cited concerns were addressed and corrected for. Popular fiction may play a role in this hesitation, with media such as *Black Mirror* and *The Matrix* that cast virtual reality in a negative light often being cited in these discussions.

Nozick's experience machine is relevant here, as Nozick (1974, pg. 43-44) hypothesised that in general people would not plug themselves into said machine. Multiple experiments suggest this to be true, with 71% to 84% of individuals opting not to plug into an experience machine (Weijers 2014; Douven & Hindriks 2018). However, there is reason to suspect that this instinctive hesitation is not necessarily tied to the virtual. Dan Weijers (2013) describes an alternative scenario named *Trip to Reality*, initially conceived of by Adam Kolber (1994). In this scenario, your current life is discovered to be the result of an experience machine, and you are given the choice to return to your more monotonous real life.

Felipe de Brigard (2010) found that when presented with this scenario, only 54% of individuals opted to exit the experience machine; a much slimmer majority than the aforementioned results. Weijers notes that similar thought experiments tended to result in people overwhelmingly choosing to remain in an experience machine. Even if one's real life was that of a "multimillionaire artist living in Monaco" (Brigard, 2010, pg. 5), those surveyed were evenly split between choosing that reality and their current life in an experience machine. Both Weijers (2013) and Brigard (2010) argue that if people are choosing not to plug into an experience machine due to an innate preference for reality over simulation, then the results of the *Trip to Reality* thought experiment do not make sense. Instead, they posit that a common cause of these differing results is a preference towards the status quo. When the status quo is not being plugged into the machine, the lack of familiarity puts us off plugging in, and when one's life has always been in an experience machine, the familiarity we have with our lives compels us to return to it (Weijers 2013; Brigard 2010).

Over the course of my research, I spoke to several individuals about the prospect of living in a reality machine, and some of them were in fact receptive to the idea. These particular individuals were overwhelmingly young adults who had experience playing in virtual reality. Of those that had those qualities, only one of them expressed persistent hesitation towards living in a reality machine. In comparison, most of the individuals I spoke to who were similarly hesitant had less familiarity with VR and had spent a greater portion of their lives without it, along with other technologies. These conversations may illustrate a broader trend of those growing up with more modern computer

experiences, and later VR, perceiving those systems as more familiar, thereby associating them with the status quo. This could make a shift to permanent virtual living seem like less of a drastic change to them, compared to readers of Nozick's literature during the 1970s.

Unfortunately, none of the empirical studies on thought experiments such as a Trip to Reality or even the Experience Machine have delineated answers based on age groups, which would be monumental in making a case for such a trend. However, Weijers (personal correspondence 2024) has noted that when surveying students, "close to half [...] indicated that they would connect to the machine". He too suggests that this is due to a growing acceptance of virtual reality and other technologies in younger generations (Weijers personal correspondence 2024).

Additionally, some parallel trends have been found that could correspond to a shift in perspective towards virtual living. In a survey of over 2000 Americans, 50% of Millennial and Gen Z individuals believed that "online experiences are meaningful replacements for in-person experiences" (Westcott et. al. 2023, n.p.), while only 19% of those in previous generations felt the same. The same survey also found that a higher percentage of the younger cohort socialised primarily online or in video games, and a majority of younger gamers found that video games supported certain emotional needs, such as confidence and self-expression. This connection between wellbeing and gaming in particular has a plethora of evidence to support it, even if particular practises by game publishers or individuals may be detrimental to wellbeing in some cases (Weijers & Munn 2024; Johannes et al. 2021; Halbrook et al. 2019).

Most significantly for our analysis, Millennial and Gen Z individuals, compared to older generations, were consistently more open towards regular usage of VR for a wide range of activities, including travel, social gatherings, education, and shopping (Auxler & Arbanas 2023). This, along with the data regarding online activity, paints a picture of younger generations as more familiar and comfortable with emerging technologies. This familiarity makes their increased usage seem more reasonable, potentially increasing their openness to permanent virtual living. Of course, this is far from definitive proof of such a trend. An empirical study on how different age groups respond to relevant thought experiments and their openness towards a virtual society project would offer much needed clarity in this area.

Regardless, even if status quo bias or any other factor has nothing to do with how people respond to questions of virtual reality, it remains an open question as to whether the hesitancy surrounding virtual reality is an immutable preference or one that will shift over time. If virtual reality does become more and more commonplace, the gut-feeling problem could very well become less and less of an issue.

## **Summary of Section 1**

By addressing these often cited concerns; deception, authenticity, placehood, and a connection with reality, we have established that many of the prudential contributors to well-being that we value about non-virtual reality can persist and thrive in a virtual society, despite potential intuitions to the contrary. In addition, these negative intuitions towards a wholly virtual life may not be inherently linked to the virtuality of the space itself, and may be diminishing as technological familiarity increases.

## **Section 2 – “Too-Goodism”**

Let us begin this section by imagining an extremely optimistic version of the virtual society; what we will call the “virtual utopia”. The centrepieces of this society are large, secure buildings that maintain the physical bodies of citizens whilst their minds experience a shared virtual world. Every citizen is aware of this arrangement, and they are free to enter and exit at any time, the virtual utopia also providing several physical communities for people to live in. The virtual world itself is abundant in landscapes and options, and provides infinitely malleable and permanent living spaces to every citizen, as explored in the previous section. In short, everyone has access to any experience they desire. Additionally, the government of the virtual utopia is transparent and accountable, and various mechanisms ensure that oppressive regimes cannot arise. The virtual utopia has no external threats, and works to ensure its long-term sustainability.

This virtual utopia has no exact analogues in popular media, at least as far as I could find. However, one could imagine it as a loose mixture of a Matrix-esque human storage system, combined with the holodecks and post-scarcity society of Star Trek (Forte 2018). In this case, the allusions to the Matrix are primarily visual. In the virtual utopia, the storage system is built by humanity for its own benefit, and is not forced upon anyone. This is in contrast to the Matrix, where individuals are trapped in a simulation that they have no control over. The parallels with Star Trek are more compelling, and lead us to the heart of this section’s argument. In one episode of the Next Generation series, entitled “The Neutral Zone”, frozen humans from the 20th century are revived by the Enterprise crew (Hurley & Conway 1988). Upon learning about society in the 24th century and the fact that “Material needs no longer exist”, one of them asks “Then what’s the challenge?”.

For most people, regularly working and acquiring resources are vital for survival. If efforts are not made to do so, you risk social ostracization, lack of food and shelter, or even an early death. Even when you successfully carry out such tasks, there are still many threats to your survival, including car accidents, wild animals, fatigue, and malicious individuals. This dynamic is altered in Star Trek, and even more significantly in a virtual utopia. Your survival is no longer dependent on the amount of work you do, and your real body is kept safe from almost any external threat. In such a future, humans may even be immortal, able to live in the virtual utopia for as long as they wish.

This scenario might invoke in you the position of “too-goodism”. This position stipulates that it is not good for one’s well-being to live a life without significant risks and with a near-limitless abundance of experiences. Jonathan Haidt (2024) claims that a similar phenomenon affects the youth of today, where insufficient exposure to real world risk, combined with unrestricted access to the internet, leads to an “anxious generation”. Similar reasoning is often used in media to counter visions of the future similar to the virtual society. The trope of a place or a life that is “too good” is highly prevalent in fiction, with the perfect world either leading characters into depression and boredom (Finkielstein 2016), or harbouring some underlying negative aspect, invoking the ‘gut-feeling problem’ that we explored earlier. Examples of the former include the Axiom in WALL-E by Pixar and Andrew Stanton (2008), the world of 17776 by Jon Bois (Tale Foundry 2024), and the Good Place in the eponymous TV series by Michael Schur (2016-2020). In both the Axiom and the Good Place, individuals are given abundant access to a wide variety of experiences, yet many of them are either bored or completely desensitized to them. Jon Bois’ 17776 depicts a world of immortal human beings that extend sporting events to last absurdly long, with a single football pitch covering thousands of square miles, all to distract themselves from the routine of their immortality for as long as possible.

This prevalence of too-goodism across media suggests that it is a common attitude, hence why it is worth addressing here. Unlike the other concerns of wellbeing that we have explored, too-goodism can be easily applied beyond virtual societies to critique all sorts of utopian futures, and perhaps even welfare systems like universal basic income. In this section, I will examine the arguments for this position, and alleviate each concern in the context of a virtual utopia.

## **An Unearned Reward**

In our current society, wealth and success carries with it a level of rarity, emblematic of someone unusually lucky, innovative, or hardworking. A successful individual able to accrue significant wealth may derive a good portion of their happiness not just from the wealth itself, but because they feel that they have earned it. More pessimistically, that enjoyment may arise from a sense of power or 'superiority' over those less wealthy than themselves, or the fact that they have a level of wealth which most people do not. Regardless of the exact reasons, some might argue that in the virtual society where everyone can experience almost anything, people will feel that they have not earned such a blissful existence, or that hard work and innovation will not provide enough benefits beyond what they already have.

We will set aside the more pessimistic possibilities for the moment, as we will go into further detail about how a stable virtual society ought to approach the hoarding of power in the final section. Let us assume, then, that the feeling of having 'earned it' does indeed constitute a large part of the happiness that wealth causes. Even if this is the case, there are means to apply that feeling in the virtual utopia. In terms of earning the abundance, I believe that all people born within the virtual utopia will have to earn that abundance by growing, learning, and maturing. I will explore this facet of the utopia further in the next section about Haidt's 'Anxious Generation' (2024). In addition, work and innovation can still provide social capital, in the form of trust and familiarity. The more involved you are with others and your community, the more your voice will be heard and appreciated. Additionally, working to improve your skills, be they technical, scientific, or otherwise, can earn you a career. Lastly, and most significantly, the malleability of each person's 'mini-universe' means that someone could reward themselves with better and better accommodations the more effort they put in each day. Most people probably would not do this, but for those who truly believe in happiness as something that needs to be earned, that option remains available to them. Whether they choose to jump right into the virtual utopia with a beautiful abode, or program their mini-universe such that they need to earn that abode through years of work, the virtual utopia is flexible enough to accommodate both lifestyles, along with many more.

## **Boredom and Monotony**

Monotony is a reasonable concern when considering a utopian future. Finkielstein (2016, pg. 106) establishes that many prominent literary utopias "frequently come to be experienced as "boring" or "unsatisfying" for their inhabitants". There are a number of reasons he points to. The first is that the very nature of an 'ideal' utopia signifies a degree of unchangeability. Any change to that ideal utopian system would only make it less perfect, since "the "ideal" is complete in and of itself" (Finkielstein 2016, pg. 106). Finkielstein notes this to be true of several utopian models, including "Plato's State, More's Utopia, Campanella's The City of the Sun, or Bacon's New Atlantis". He argues that in these utopias, the unchangeable social systems and predictable routines of their inhabitants would foster boredom.

Now, if we are to design a utopia that avoids this issue, we should not allow for every aspect of it to be easily changeable. A government that is able to extend their time in office indefinitely and pass

laws letting them kill anyone they wish without resistance would be a nightmarish scenario. Instead, a utopia looking to prevent monotony should be 'socially and environmentally dynamic', and to an extent, should be independent from the daily lives of its citizens. In this sense, 'independent' describes a relaxing of specific social expectations and classes imposed by the state, as is the case in both Utopia and State (Finkielsztein 2016), in favour of fewer and more fundamental principles of conduct. Such principles might include the forbidding of murder, corruption, and impersonation. More specific principles of sociality, such as etiquette, slang, entertainment, and manners of celebration, would be allowed to vary between different communities across the nation. Every citizen would be guaranteed the freedom to associate with as many or as few of those communities as they wish.

Under this framework, I will define 'socially dynamic' as the descriptor for a society that allows individuals to decide how they socialise and who they socialise with, and allows those decisions to change over time. In short, a socially dynamic society has a malleable social environment, where communities themselves are free to morph, grow, fade away, and be created. An 'environmentally dynamic' society is one in which the 'physical' make-up of any given community; such as buildings, pathways, and trees, can be significantly altered over time. Lastly, the malleability of both the social and physical aspects of any given community should be primarily the result of the community itself.

There are a few reasons why independent and dynamic communities in a utopia can ease boredom. Having access to a wide range of different communities can make socialisation less predictable, and make travelling throughout the utopia more exciting, since each location has a distinct character to it. Because communities can change over time, the future becomes somewhat more uncertain and therefore more interesting. People will be more incentivised to participate in their evolving communities, either to make the changes they want or to prevent changes they do not want. Lastly, people will feel more comfortable expressing themselves in their own unique way without feeling pressure to conform from an overarching state-enforced social system. If someone's community is less inclined to a particular expression of identity, that person can move to a more accepting place, or build a new community with like-minded individuals. This is particularly salient in the virtual utopia, since space, distance, and resources are not significant hurdles to community switching or building as they are in the real world.

The independent and dynamic utopia (or 'IDU') also addresses other pitfalls that Finkielsztein brings up. He states that "Boredom is commonly linked to a paucity of control, an excess of freedom [...] or to [a significant] lack of freedom" (Finkielsztein 2016, pg. 108). The IDU allows individuals substantial control and purpose within their communities without allowing them to be so free as to be detached from those communities. The aim of the IDU is to find the happy medium between feelings of powerlessness due to excessive control and 'anomie'; the "state of lawlessness [that] often evoke[s] a correlative sense of insecurity and anxiety" (Finkielsztein 2016, pg. 108). Examples of the latter include political and economic instability. While such instability can be interesting and exciting to a certain degree, they can have highly negative consequences not only to mental health, but to many other aspects of life beyond that. As a result, a good IDU will work to prevent excessive instability where it produces more harm than good.

Assuming our virtual utopia is also an IDU, we still need to consider the day to day lives of citizens in this dynamic society; specifically the level of satisfaction possible at the individual level. At the beginning of this section, I noted the discrepancy in the amount of effort needed for most people to survive in the real world compared to the virtual utopia. I will call this effort 'necessary effort', as it is necessary for survival. For some, a life without necessary effort might seem lesser than one with it; without the struggle to survive, they might say, life becomes meaningless. This is reflected in the

notion that life requires risk to be meaningful, such as the risks of suffering, loss, and death. Necessary effort is satisfying not only due to the rewards it might bring, but also because it temporarily moves you further away from the possibility of death and suffering. Therefore, some might argue, there is a unique satisfaction that comes from work which ensures your survival and the survival of those around you, compared to work that simply fulfills personal, non-essential desires. This idea is not unreasonable. Hunting a dangerous beast that threatens your tribe and can provide enough food to feed them for weeks would likely be far more satisfying and meaningful to you than hunting that same beast simply out of boredom.

Of course, the vast majority of humans do not hunt game to survive anymore. The nature of necessary effort has changed substantially. Most people around the world work in a wide variety of jobs in exchange for currency, allowing them to purchase or rent essential items for their survival, such as food, clothing, and shelter. As a result, beyond the payments you receive, the job you have and the amount of work you put into it is not guaranteed to provide as clear and tangible benefits to you and those around you. If this is the case for you, then the necessary effort you put in may feel less satisfying. Returning to our example of hunting a beast, the correlation between the work of hunting and the reward of food and safety is much clearer than between the work of an office job and the reward of being able to purchase food from elsewhere.

Let us assume, then, that the argument for necessary effort is true. Given that the nature of necessary effort has changed a lot already, perhaps we can modify it once again; not just to fit into our virtual utopia, but to give us a reasonably strong correlation between work and reward.

An important component of necessary effort is 'social effort'; devoting time and energy into creating and maintaining relationships. The importance of this practice is very clear in a hunter-gatherer society, as not getting along with your tribe makes them less likely to share food and water with you, or help you recover from an injury. But there's a second component to social effort that applies even in a post-scarcity utopia, and that is its vital positive impact on mental wellbeing.

Humans are social beings, and a lack of sufficient social interaction is known to have a severe negative impact on mental health (CDC 2024). As a result, there is an innate incentive to maintain one's social wellbeing, to get along and connect emotionally with other people. Maintaining relationships with people, even enjoyable ones, requires time and energy, and this is especially true when it comes to romantic relationships and raising children. The social incentive would still persist even in a society where all other material needs are met, assuming that we are consistent with our idea from Section 1 of limiting AI companionship. We intuitively understand, for example, that a billionaire with no social connections would most likely be miserable, even though they are easily able to live in luxury for the rest of their days. Therefore, individuals in the virtual utopia would still spend social effort, due to the value they place in their friendships, partners, and children. After all, relationships would be one of the rare things in the virtual society that someone could lose. Since 'social effort' is a form of 'necessary effort', necessary effort would therefore remain as a requirement for living well. In addition, the social rewards would have a strong link to the social effort, given that the same people are involved.

One possible objection to this reasoning could be that social effort is not enough to constitute a satisfying amount of necessary effort. People, after all, "want to do certain things" (Chalmers 2022, pg. 305). More specifically, the argument may go, people deeply desire to be productive and valuable. If social productivity and value is all that remains in the virtual utopia, and if such things are insufficiently satisfying, then that desire of individuals to be productive and valuable will not be satisfied. If we assume it to be true that social effort is insufficient for wellbeing, and a greater

degree of productivity is needed, then we must show that there is meaningful productivity to be found in the virtual utopia that extends beyond the social.

Recently, I lent a hand to my brother by setting up some furniture for his new business. Even though this act may have been somewhat beneficial to me socially, in that it could have contributed to my brother having a positive opinion of me, and it felt nice to spend time with him and help him, the social benefit was not my only source of satisfaction. The act of putting together the furniture, and the knowledge that I was contributing to a larger, tangible project that could have an impact on several other people across time, was also satisfying to me. What is most interesting here is that my brother's project of starting his own business was not obviously necessary for his survival. He has the experience and good standing necessary to easily be hired for many jobs. In fact, given his investment in his business, it may prove detrimental to his survival-related resources if he is unsuccessful. It is a project driven purely by a desire to live a certain kind of lifestyle and do the sort of work that brings him joy and gives his life more meaning.

These acts of building furniture and starting businesses are examples of fulfilling productivity. Fulfilling productivity is a form of necessary effort in the same way that social effort is; they both have benefits to mental health and wellbeing. Occhipinti et. al. (2023, pg. 3) state that "Being productive is important for mental health. Studies dating back to the 1930s highlight [...] both the positive impacts of being engaged in productive work and the negative effects of unemployment". Caitlin Nevins (McClellan Hospital 2024, n.p.) notes that "We tend to value those days when we get a lot done," and that "A sense of mastery and accomplishment is good for our mood". If this is true for you, then it would remain so even in a virtual utopia, where productivity is not necessary to acquire basic resources or shelter. Fulfillment born of effort could still remain an important part of your everyday life.

To illustrate this, let us imagine that you live in a virtual utopia, and you have a dream of running a cozy coffee shop. You work with locals in your community to find a space that suits you, spend time determining how it will look both inside and out, and ultimately decide to set up the interior space yourself, rather than automating the process and sidestepping the hard work. This is because you enjoy painting, carpentry, and arranging things in a tactile manner, and the virtual utopia gives you the freedom to choose to do so. Eventually, you finish setting up your coffee shop, and now you can get started on the kind of work that fulfills you. The shop does not even need to be popular, since money and resources are not an issue in the virtual utopia. You can simply spend as many days as you like, being productive and bolstering the community in a way that suits you.

Just like in the example of my brother starting a business in the physical world, the same sort of indirect necessity is apparent here. You do not need to run a virtual coffee shop to survive, just as my brother did not need to start a business to survive. Both scenarios are instead allowing for meaningful productivity and the fulfilment of ambitions. The only difference between the two examples is the medium; the physical world and the virtual one.

One argument against this perspective is that real coffee shops have a tangible benefit to their communities. They can provide important nourishment and caffeine, while a virtual coffee shop only gives the illusion of those things. If there is no real benefit to a virtual coffee shop, that might put a damper on how fulfilling it can be for those working there. While it is true that virtual coffee shops do not provide genuine nutrition, virtual coffee shops are still beneficial to virtual communities. They offer a place for people to enjoy the sensation of food and drink, and to talk with one another, as well as meet new people. Those visiting from other parts of the virtual utopia can sit down with a coffee and plan out the rest of their trip. The more shared activities there are available to a community, such as coffee shops, the more that community can develop a sense of cohesion and

identity, thereby helping to prevent isolation. As a result, running your own virtual coffee shop can still be both fulfilling and meaningful.

Additionally, there are several advantages that a virtual utopia has over reality when it comes to fulfilling work. As mentioned earlier, you could continue running a coffee shop even if only a handful of people use it, all without worrying about your bottom line. Unlike reality, there are no limitations on the kinds of work you can do, nor do you need to worry about making ends meet. You can incorporate that aspect if you wish, but it is not essential, and if you do incorporate it, you would not face potential poverty or homelessness if your business fails. Your fulfilling work could range from building model ships to power-washing spaceships, from writing books to climbing mountains, to advising a king in a medieval community. And if you ever need to take a break from your work, you can easily do so risk-free.

Despite all of that, some people may still want a career in the virtual utopia that better resembles those of the modern era; jobs that are vital to keep things running and carry substantial consequences if done poorly. Fortunately for them, even a virtual utopia carries with it several highly important responsibilities across a wide range of disciplines. These include jobs that require citizens to perform tasks in the real world, such as repairs, maintenance, cleaning, and logistical oversight, as well as jobs that can be carried out in the virtual utopia, such as governance, scientific analysis, design, and advisory roles. Many of these positions and their structures will be discussed in greater detail in the next section, but their inclusion here is to illustrate that such jobs will still remain necessary and available in the virtual utopia to those who desire them.

Overall, there are many reasons to believe that citizens of a virtual utopia are not doomed to inevitable boredom. Individuals can meaningfully participate in communities that are socially and environmentally dynamic, and putting in effort will still be required for both social acceptance and mental health, just as it is in our current reality.

## **Overabundance**

Another reasonable concern of bountiful utopias is their ability to provide limitless experiences, both in terms of quantity and variety. While this facet helps alleviate boredom, one concern is that individuals exposed to this will quickly overindulge in the most satisfying experiences possible, rapidly becoming desensitized to everything the utopia has to offer. Alternatively, the sheer number of possible experiences and directions one could pursue in the utopia might lead other people to inaction through decision paralysis, where even choosing what to do each day could become a stressful exercise, like picking one of Netflix's 5000+ movies to watch (JustWatch n.d.).

This is a valid concern, but not one unique to the virtual utopia. Overindulgence amongst those with access to significant resources is a common trope. Finkielstein (2016, pg. 108) describes the melancholy of the "artists, aristocrats and bourgeoisie in nineteenth century Germany" brought about by excess. Blaise Pascal (1887) notes that a king would find himself unhappy if he were left alone without the distractions of other people. Given the associations this overindulgence has with wealth and power, it is worth examining the overall life satisfaction of those with plenty of it, as they are the most similar real-world example to our citizens of the virtual utopia. If there is a way to have abundant access to pleasurable experiences and still maintain a comparably happy life, either through moderation or some other practice, then people can do the same in the virtual utopia.

One complication with this approach is what life we choose to compare the wealthy life to. There are various confounding factors at play when distinguishing, say, the happiness of a wealthy person and the happiness of an impoverished person, beyond increased access to experiences. Wealthy

people are less likely to deal with the stresses of loan repayments, or work multiple jobs with long hours, or develop illnesses due to poor insulation or inadequate health insurance. To ensure that these results best reflect a difference in access to experiences, as would be the case when comparing someone living today with a citizen in the virtual utopia, we need to mitigate such factors.

One method is to instead differentiate between those with enough money to live comfortably-- what many would call the 'middle class'--and those with enough money to buy almost any experience they might desire. This way, the focus is tied more-so to the variety of experiences possible. To further reduce confounding factors, it is also worth looking at the happiness over time of those who have experienced a less affluent lifestyle before transitioning into a wealthy one, such as lottery winners. Such a transition will show the relative happiness of the same person before and after gaining access to an abundance of experiences, similar to seeing the same person outside and inside of a virtual utopia.

These will be my metrics; the relative happiness of wealthy and 'middle class' individuals, and the relative happiness of individuals before and after receiving a substantial fortune. Measuring happiness objectively is difficult; Frances Stewart (2014) notes that the manner, phrasing, and philosophical preconceptions used when asking individuals about how happy they are can significantly impact the results. To keep things simple, we will look at signs of unhappiness; i.e. rates of depression and suicide, to start off with.

McMillan et. al. (2010, n.p.) note an "inverse association between income and psychological distress", and the top 25% of earners not showing an increase in such distress compared to the second and third highest quarters of earners. Kawachi et al. (2010) and Patel et. al. (2018) both note evidence across several studies that an increase in income corresponds with an increase in mental health. In addition, Ettman et. al. (2022) found that, in a majority of studies, wealth in terms of assets also had a negative correlation with depression. In terms of comparing middle-income individuals with the wealthy, Prins et. al. (2015) found that individuals making over 120,000 USD a year showed consistently lower rates of depression or anxiety compared to those making between 40,000 and 60,000 USD, and largely similar rates to those making between 60,000 and 100,000 USD.

Additionally, a study by Lindqvist et. al. (2020, pg. 2703) indicated a sustained increase in life satisfaction experienced by winners of substantial lotteries in Sweden. They noted that "The effect shows no evidence of fading over the time horizon for which we have data and is robustly discernible over a decade after the lottery event". In a study from decades prior, the lottery winners involved reported similar levels of happiness to the control group (Philip Brickman and Dan Coates 1978). While this latter study highlights an eventual return to an average level of happiness following the high of the lottery win, neither of the studies indicate a decline in happiness below the average.

What these studies all seem to indicate is that high levels of wealth do not correlate with negative mental health effects, nor a decrease in wellbeing below the average. A recent meta-analysis of socio-economic status and subjective well being affirms this further, with one of its conclusions being that "[...] money and resources, whether objectively reported or subjectively perceived, is significantly linked to SWB [subjective well being]". (Tan et. al. 2020, pg. 1011). Lastly, a 2025 meta-analysis (Oh 2025, pg. 20) noted the following:

"Practically, the findings suggest that financial satisfaction could indeed play a strong explanatory role in immediate well-being—across multiple indices of well-being including life satisfaction, emotions, social satisfaction, and health, financial satisfaction had the strongest and most consistent relationship with well-being in the short term. Moreover, increases in

financial satisfaction appeared to occur alongside corresponding increases in well-being over time, and these effect sizes were large. [...] Those who earned more income were moderately more likely to experience improvements in well-being or to experience weaker declines in well-being. [...] The results suggest that both [money and financial satisfaction] play unique explanatory roles for well-being.”

Overall, there seems to be a lack of any notable reduction in mental health, happiness, or well-being of wealthy individuals on average compared to other income groups, and even evidence suggesting the opposite is the case. This indicates that the overabundance of experiences they have access to because of that wealth does not negatively impact their wellbeing. There is little reason, therefore, to suspect that overabundance would negatively impact citizens of the virtual utopia either.

In fact, the virtual utopia may be capable of curbing a particular negative aspect of overabundance. This aspect of overabundance I wish to explore is decision paralysis. Barry Schwartz (2015, pg. 122) argues that an overabundance of choice creates a “paralysis”, a “misery-inducing tyranny”. He cites examples of studies where individuals are less likely to make a purchase or feel satisfied with the choices on offer, when there is a large number of options to choose from, instead of a smaller number of options. Based on these and many other cited studies, Schwartz (2015, pg. 123) states “it now seems clear that under a broad range of circumstances, people find a large number of options paralyzing rather than liberating”.

When exploring this phenomenon further, Schwartz (2015, pg. 123) invites us to consider ‘maximisers’; those more likely to “seek the very best option available in a wide range of choice domains”, and ‘satisficers’; who are more likely to “choose the first option that surpasses some absolute threshold of acceptability”. He notes that while satisficers may not be significantly impacted by a greater abundance of choices, maximisers may struggle as the number of options increases, because “as the number of choices in a domain increases, so too does the cognitive work required to compare various options, along with the possibility of making a “wrong” or suboptimal choice” (Schwartz 2015, pg. 123). For example, a maximiser trying to pick a board game for her friends would have a much easier time finding the one game that her friends would enjoy the most when there are only five board games on the shelf, compared to fifty. Meanwhile for the satisficer, the difference between five or fifty boardgames is not much cause for concern. If anything, more options could mean a higher likelihood of finding a game that meets their standard of quality.

Of course, it is worth noting that most people would likely use both methods throughout their lives, depending on their mood and the sort of choice they have to make. It would be reductive to define someone as ‘solely a maximiser’ or ‘solely a satisficer’, even if they strongly prefer one over another. Regardless, there are several factors that Schwartz (2015, pg. 126) believes “undermine the objective benefits that ought to come with increased choice”. These factors affect both groups, but particularly maximisers, and include regret or “buyers remorse”, “missed opportunities”, and “social comparison”.

In the virtual utopia, many of these factors are either reduced or non-existent. Buyers remorse, a potentially large driver of maximising behaviour to begin with (Schwartz 2015), is much less relevant in the day to day activities of its citizens, since many objects and experiences do not require resources or effort to obtain, and can easily be swapped out if they do not suit. Missing out on one experience over another is not an issue. Even your lifespan could be extended to give you enough time to experience as much as you wish before moving on.

According to Schwartz (2015), social comparison’s impact on decision paralysis relates to the observation of the successes of others, and a desire to make good decisions to reach that same level

of success; be it financial, social, or another factor. This naturally becomes more difficult to achieve the more choices you have to make. The virtual utopia is egalitarian as a whole, incentivising the pursuit of your unique happiness and spending quality time with those you are close with. Financial and social mobility are irrelevant to maximising one's well-being, though it is possible that in the virtual utopia, the pursuit of happiness itself could become a target of social comparison. An example of this can be found in the Netflix series *Carol and The End of The World* (2023), where large amounts of the population are all trying to squeeze as much fulfillment as possible out of their final months before the world's destruction.

Schwartz (2015) offers a few suggestions to counter the negative aspects that unrestricted freedom of choice can bring. As mentioned earlier, he notes that maximising strategies exacerbate decision paralysis, and postulates that "the cultural pressures in a post-industrial capitalist society [...] might lead to the development of maximizing tendencies" (Schwartz 2015, pg. 134). If this is true, then our virtual utopia may consist largely of satisficers, since there is no pressure either personally or socially to accumulate resources or wealth. However, the primary suggestion Schwartz offers is that a society could "sensitize people to the costs of choice and the benefits of constraints. People who appreciate the benefits of constraints might seek and embrace the constraints that arise out of membership in a close social network" (Schwartz 2015, pg. 134). This approach may seem contrary to the abundant nature of the virtual utopia, but in fact, this is how many of its virtual communities could run. Communities in the utopia are intended to give that very same sort of structure to the lives of those who live in them, and different communities will offer different lifestyles, some more or less flexible than others. The digital nature of the utopia can aid in voluntarily limiting choices as well. An AI agent, upon request, could help you narrow down your options, or even surprise you with an option it thinks you would enjoy the most, if you wished. Such services would always be available to use or to not use; for as often, or as little, as you wished.

## **Haidt's Anxious Generation**

As noted earlier, Haidt (2024) has expressed some concerns about how children are being raised in the current era, particularly with how it is impacting their development and adulthood. This approach is a confluence of "two trends - overprotection in the real world and underprotection in the virtual world" (Haidt 2024, pg. 5), which, according to him, runs counter to how children ought to develop, leading to a significant increase in mental disorders in Gen Z adults (Haidt 2024).

Whether this claim is true or not is largely irrelevant; we can even be charitable and assume that it is. What is notable is the prevalence of sentiments similar to this across the population. Mukherjee (2021, pg. 646) notes:

"Claims that digital media is turning children into 'digital zombies' and 'psychotic junkies' are ubiquitous in the news media [...] These sensationalised media accounts have linked children's ICT [Information and Communication Technologies] use to the decline in their outdoor play as well to their physical and mental wellbeing."

According to a PewResearch report, "71% of parents of a child under the age of 12 [...] are at least somewhat concerned their child might [...] spend too much time in front of screens, including 31% who are very concerned" (Auxier et. al. 2020, n.p.). Another report from The Wildlife Trusts (2021, n.p.) claims that 75% of adults "believe children do not spend enough time outdoors". Even if we apply a generous margin of error to these results, a significant portion of adults seem concerned about these factors.

Of course, the amount of time spent in front of screens or outdoors does not directly translate into overprotection or underprotection from either. A child could spend hours and hours outside each day and still be overprotected from any harmful aspects. But this type of vigilance would likely become more challenging the more time the child spends outdoors, as the odds increase of the parent becoming exhausted, distracted, or otherwise unavailable. This means that the amount of time a child spends outdoors does have some impact on how protected they are from it; their 'exposure' to its more dangerous elements.

This highlights some overlap between the wider concerns of adults and the concerns of Haidt. The concern of how unexposed modern children are to the outdoors also elucidates the underlying concern; children are less free to 'make mistakes' and learn from them than they once were. Common sentiments that reflect this concern include those of Mark Gregston (2020, n.p.), who describes how the "unstructured, self-discovery play" of his youth was an "essential rite of passage that built self-confidence and courage in children". According to a Mott poll (Child Health Evaluation and Research Centre 2023), most parents of children aged 5 to 8 and 9 to 11 seem to share the sentiment that unsupervised play is important. This is despite the same poll showing higher amounts of protective tendencies from those parents than their sentiments would suggest.

The question remains on whether unrestricted internet play can replace unrestricted outdoor play. Here is where Haidt's argument comes back in. He believes that the current internet landscape is unsuitable for normal childhood development, stating that "Virtual interactions with peers do not fully compensate for [the] experiential losses [of outdoor play]" (Haidt 2024, pg. 8). In addition, Haidt notes that:

"Those whose playtime and social lives have moved online [find] themselves increasingly wandering through adult spaces, consuming adult content, and interacting with adults in ways that are often harmful to minors" (Haidt 2024, pg. 8).

Let us assume that Haidt and his peers are correct, that children are overprotected in reality and underprotected in the digital realm. It seems at first that the virtual utopia would make this problem worse. Children would be isolated from the outdoors almost entirely, all the while free to explore a consequence-free digital space without limit. However, I believe that there is a means for individuals to have an even better childhood in a virtual utopia than in our current reality, both in terms of fulfillment and proper development.

In the previous section, I introduced the concept of 'maturity worlds' where children can spend time with their caregivers, make friends, and have the freedom to grow in an environment that provides challenges and risks without the threat of severe long-term harm. It is this concept that I wish to discuss further, as I believe it best addresses the concerns of Haidt and his peers. The creation of spaces designed for childhood development is nothing new. Setting aside schools and other educational institutes, the promotion of playgrounds and natural outdoor spaces for children can be found in the works of nineteenth century authors such as Fröbel (1886) and Carlyle (2008). And just as a playground can steer children away from dangerous motorways, a maturity world can shelter children from the wider digital utopia until they have the temperance and experience to live in it healthily. To assuage the worries of Gregston (2020, n.p.) regarding "ubiquitous government-regulated, sterile pre-fab playgrounds", the maturity worlds could more so resemble the 'riskier' play spaces of his youth, without being truly deadly or risking life-altering injury. Thus a happy medium is reached, where play and experimentation can still come with the risk of experiencing scrapes or broken bones. If such experiences are the best way for children to develop, then maturity worlds can accomplish this.

In order for maturity worlds to be a viable solution to ‘anxious generations’, they must stand up to scrutiny, and there is much to scrutinise about them. First of all, it may seem immoral to confine children to an imperfect world when they could be living in a utopian one. This objection is particularly salient since learning experiences involving injury are an intended part of the maturity world, while pain is entirely optional in the virtual utopia. Yet without such experiences and the delayed access to utopia, the maturity worlds cannot fulfil their purpose of preventing ‘anxious generations’ by allowing children to make mistakes and learn from them. While a parent might feel some guilt about denying their child access to the virtual utopia, the same is true when a parent denies their child a diet of sweets, or drops them off at kindergarten for the first time.

In the previous section, I raised the issue of whether or not children in maturity worlds are informed of the virtual utopia. I described a few options, and the pros and cons of each one. What I did not mention is that there may be an alternative means of letting children know about their future in the utopia without it compromising their drive to grow. The requirements for leaving a maturity world could be based upon one’s emotional and mental age rather than an arbitrary numerical age. This may be determined by a decision from the primary caregivers, the adolescent making a case for themselves, psychological evaluations, a challenging rite of passage, or some combination of the four. By informing children that the wider virtual world is there for them when they are ready, they are more likely to foster the growth necessary to reach it.

Another potential concern is that if children all grow up in maturity worlds, they might be very similar to one another, lacking in diverse backgrounds or perspectives. This does not have to be the case, however. Different maturity worlds could incorporate different climates, different architectures, different cultures, and different histories. Additionally, every child would have their own caregivers, and different children that they share the maturity world with to influence their development. The one perspective that may be lacking is an impoverished one, since caregivers would want a reasonable standard of living in maturity worlds. Even so, empathy for those who are worse off can still be taught, either through learning about history, witnessing relative impoverishment, or even through their own personal struggles. The potential loss of a particular perspective does not justify leaving children in poverty, either in the real world or in a maturity world.

Let us analyse now the advantages of a maturity world compared to modern childhoods, in the context of Haidt’s concerns. Firstly, caregivers can be more confident in letting their child explore and develop, since nothing in a maturity world can cause severe lasting harm. Unlike in the real world, where some children are faced with enormous amounts of difficulty and others faced with very little, maturity worlds ensure a sufficient yet not overwhelming amount of difficulty for all. If a child’s caregivers are too coddling, educational institutions in the maturity world can push that child a little further, and if a child’s caregivers are too harsh, those same institutions can be more lenient. These schools can function indefinitely without fear of budgetary pressures or a lack of staff. While real-world schools are limited in what they can teach, maturity world schools would have the flexibility to provide learning opportunities for nearly any interest or passion a student could have.

Lastly, if technologies equivalent to computers and smartphones are available in maturity worlds, they can easily be made safer and more limited. Perhaps the simulated internet could lead teenagers to healthy hobby groups within the maturity box, or video games could only work for a few hours each day. Even if we are not comfortable with restrictions like these, the internet of the maturity world could still filter out the harmful adult spaces and ‘addictive content’ that Haidt (2024) is concerned with. In particular, given his concern on the matter, there are means to make social media less addictive as well.

One study explored the impacts of a custom-made browser extension called NUDGE which requires users to perform an increasingly difficult action each time they try to access Facebook (Barclay, Holzer, & Purohit 2020). Alongside other habit breaking tools that NUDGE employed, the extension was found to reduce time spent on Facebook, made that time more enjoyable, and seemed to increase how mindful users were of their time spent on social media. In the real world, applying sweeping reforms to create an internet more suitable for childhood development is more challenging. Unlike in the virtual utopia, children do not have access to a more suitable secondary internet, instead relying on the same websites and platforms that adults do. Additionally, these websites have financial incentives to engineer a user experience designed to keep those users engaged and collect their behavioural data (Zuboff 2019), often running counter to genuine ethical design philosophy (Saura et. al. 2021). Such obstacles to suitability would not present themselves in the virtual utopia.

## **Summary – What is the Challenge?**

In a world where “material needs no longer exist”, Captain Picard states that the challenge is “...to improve yourself. To enrich yourself. Enjoy it” (Hurley & Conway 1988, n.p.). The virtual utopia, much like Star Trek’s United Federation, gives its citizens the freedom, the support, and the challenges necessary to grow into strong and capable individuals, enriching themselves and the lives of those around them in the process. The concern of too-goodism, while sometimes valid, does not apply strongly in the context of ideal virtual societies and maturity worlds.

## **Section 3 – Abuses of Power**

Many who are hesitant to live in a virtual society believe that there is an unreasonable amount of trust one has to place in those running the simulation. The government of a virtual society has to consider both the internal/nonphysical aspects of that society and the physical infrastructure that the simulation is dependent on, along with the means by which physical people can interact with the simulation. Not only are there several avenues for corruption here, but the ease of harm and oppression is far greater for citizens of virtual societies than physical ones.

Dictators in the real world are ultimately bound by the laws of physics and their dependence on others to enforce their will. As CGP Grey (2016, n.p.) puts it, “No man rules alone”. This is not the case in a virtual society. If a tyrant held complete control over a simulation and all of its facets, he could create an inescapable hell of perfect agony all from the comfort of his computer. This is not to say that physical societies are free from abuse, corruption, and suffering. However, it is definitely understandable why someone would hesitate to live in a society where potential oppressors could exert cruelty instantly and inescapably.

As is the case in our current non-virtual societies, we cannot alleviate this concern entirely. But there are measures we can take to increase citizen trust and decrease the odds that a virtual hell is realized. These five overarching measures, which I will refer to as ‘Axioms of Virtual Abuse Prevention’ (or ‘AVAPs’), will be rooted in ideal practices and strategies for both virtual worlds and resilient democratic governments to maintain cohesion and happiness. Each will be used to develop a secure framework of trust and freedom for citizens of virtual societies.

In the first section, I mentioned two areas that a virtual society has to consider, namely the real world that hosts the simulation, and internal systems designed to prevent or deter abuses of power. These areas of consideration will also come into play for each AVAP.

## **AVAP 1 – Disincentivization**

First of all, it is worth asking what motivates a person to seek power and what might drive them to abuse it. Identifying these factors will be useful in formulating a society that mitigates them, so that those in leadership positions and other important roles in our virtual society are acting more as public servants rather than would-be tyrants.

Foulk et. al. (2020) outline three “moderating factors” that can increase or reduce abuses of power, the first being ‘individual differences’ Higher levels of particular traits in those wielding power, such as ‘agreeableness’ and general empathy, may reduce “the tendency for power to manifest in self-interested or antisocial behaviors” (Foulk et. al. 2020, pg. 4). In contrast, aspects of traits such as selfishness, ‘dominance orientation’, narcissism, Machiavellianism, and psychopathy, were more conducive to abuses of power. The study, in a broader sense, notes that:

“Individual differences can not only weaken the negative effects of power, but can also be associated with power being expressed in positive, prosocial ways. Many of [these] traits [...] are related to the way powerholders see the social environment and the importance they place on social relationships with others [...] They also point to the possibility that power may mean different things to different people [...] people who value others [may] experience power in a different way—one that motivates them to care for and attend to the needs of others” (Foulk et. al. 2020, pg. 4)

This leads us on to the second moderating factor, the ‘construal of power’, or “what having power means to the powerholder” (Foulk et. al. 2020, pg. 6). Conceptions of power as an ‘opportunity’ or as granting greater personal freedom were, across various studies, correlated with an increased likelihood of antisocial behavior (Foulk et. al. 2020). Meanwhile, conceptions of power as a ‘responsibility’ or one that emphasizes “the need to rely on and care for others” tended more towards positive social behavior, and “decreased self-interested behavior over others” (Foulk et. al. 2020, pg. 7).

The third moderating factor is the broader context of each situation, the facets of which include “perceived constraints on their power, the features of the task environment [...], and characteristics of others in [that] social environment” (Foulk et. al. 2020, pg. 5). ‘Perceived constraints’, such as the notion that others consider you incompetent, were found to potentially lead to antisocial behaviors, while ‘actual constraints’ such as sanctions may reduce such behaviors.

In terms of the workspace, or ‘task environment’, Foulk et. al. (2020, pg. 5) noted that “powerholders may generally be less prone to act in self-serving and agentic ways when they perceive important task interdependencies with others [...] or when the task environment itself provides opportunities for prosocial versus antisocial behavior”. Lastly, in regards to the social environment, Foulk et. al. (2020, pg. 6) posit that “when the context makes powerholders aware of others' dependence on them, or their responsibility for others, the negative behavioral tendencies associated with feeling powerful are either weakened or reversed”.

There is no reason to believe that these results are not applicable even in cases of high-level leadership. Notably, some of these findings mirror research done on world leaders by Nai and Toros (2020, pg. 1), which found autocrats “score significantly lower in agreeableness” than non-autocrats,

and also “score significantly higher than non-autocrats on the three traits of the Dark Triad, and especially psychopathy”.

Based on these findings, we can formulate a few guidelines to disincentivize abuses of power, per AVAP 1:

1. Incentivising qualified individuals who are empathetic, and value others substantially, to assume key responsibilities, and disincentivizing those with substantial narcissism, Machiavellianism, and psychopathy from assuming those responsibilities.
2. Foster a culture that considers key jobs to be responsibilities rather than opportunities, and reflect that consideration by limiting what can be gained by having said jobs.
3. Have binding standards and practices that everyone in key fields are aware of, designed to curb abuses of power.
4. Organize these key positions in such a way that fosters interdependence and reliance on everyone involved in each key field.

I will address each of these guidelines going forward, but there are a few more we can add to help further disincentivize tyranny in the virtual society. These additional guidelines will address the motivations for seizing power and doing harm.

One possibility is that the motivations of an aspiring subversive leader may be largely ideological; they push against the systems of government to seize power, harming people along the way, in order to bring about some ‘greater good’. Certain revolutions in history are framed in this way. Bristow William (2017, n.p.) states that “The political revolutions of the Enlightenment, especially the French and the American, were informed and guided to a significant extent by prior political philosophy in the period”. A less favorable example would be, for instance, someone obtaining power in order to outlaw any literature or conversation that is contrary to their ideological vision.

Another possible motivation is a desire for security; our aspiring subversive leader might sense that their current place in society is disadvantageous or even dangerous, and seek out a means of personal security and control through political power. Yet another motivating factor could be a twisted enjoyment of holding power, influence, importance, or control over the lives of others.

Lastly, one might be motivated to tyranny by hatred. This hatred could be directed at individuals in positions of power, particular sociocultural groups (discrimination), a broader hatred of people in general (misanthropy), or some combination of these. The usage of hatred as a motivating factor to bring about tyranny is not uncommon. To carry out or permit violence against a person or a group of people is indicative of a disdain or contempt for them, and hateful leaders throughout history have carried out or permitted horrific acts against people. As Rummel (2018, pg. 46) describes it, “Genocide, massacre, and human slaughter; pillage, rape, and torture have been much more common [through history] than war and revolution”.

With these motivations in mind, we can add our last guidelines to disincentivize abuses of power:

5. Enable different ideological voices to be heard without elevating ideologies that promote hatred or authoritarianism.
6. Provide safety, security, and good living standards equitably.
7. Encourage mutual respect and empathy.

The next step is to look at the virtual society and see how it can support these guidelines, both internally and externally. Fortunately, many of these guidelines can be supported through the consequences of supporting a different guideline.

A significant example of this is Guideline 6. As noted in previous sections, every person in the virtual society is given access to their own infinitely malleable personal space, while their real bodies are kept secure and healthy in non-virtual reality. If we assume that the virtual society is stable, then the conditions for Guideline 6 have been more or less achieved; every individual is safe, secure, and has access to any standard of living they desire. The fulfillment of this guideline helps in fulfilling Guideline 2, since there is little to be gained beyond the high standard of living that all citizens are entitled to. This also aids in curbing practices like bribery and corruption. As such, there is little benefit to working beyond the satisfaction of helping and supporting others, making each profession more of a responsibility than an opportunity.

This perception also supports the development of mutual respect and empathy per Guideline 7, since workers in the virtual society would perceive each other as being responsible and valuable citizens, and less so as means for greater opportunities down the road. Guideline 7 can be reinforced further by the implementation of an interdependent working environment as per Guideline 4, which would enable individuals to recognize what each of their peers brings to the table. The virtual society is better suited for creating such an environment than our non-virtual societies, since it would enable the rapid development and testing of different working environments; including workplace layouts and management structures, that best suit workers in any given scenario or project.

The perception of work as a responsibility may also serve as a soft filter of sorts for pulling kinder people into those responsibilities, as per Guideline 1. As for Guidelines 3 and 5, we can imagine penalties for actions such as intimidation in the workplace, or flagrant homophobia in public. These penalties may be lighter in accidental cases, and could include warnings, a requirement to rectify the situation, temporary confinement to one's home, or a requirement to meet with a therapist.

With all of this being said, there are still potential flaws in how the virtual society may approach each guideline. Some might argue that the opportunity to improve one's social standing in the community could be a motivation that slightly skews the framework of work as responsibility. Additionally, despite the guidelines, a social divide may form between those who choose to work and those who do not, with the latter group potentially being perceived as lazy or irresponsible. Such divisions may engender resentment between social groups rather than dampen them, possibly leading to hatred and ultimately tyranny.

The manner in which these concerns are addressed could be the focus of innumerable papers, and I do not doubt that there are solutions to them. Fortunately, even if some faults at this stage do turn out to be inevitable, there are more AVAP's that can be put in place in case this first line of defense fails.

## **AVAP 2 - Immutable Safeguards**

An essential part of any democracy is a robust series of laws and procedures to curb potential threats to the democratic process, both within political institutions and outside of them (Boese et. al. 2023). In the virtual context, such regulations can take on forms more akin to physical laws, such as needing to complete a lap in a racing game in order to score a point. In such a case, the point system is not bound by a social contract or how people might feel about the game, rather it is an intrinsic property of that virtual universe.

One crucial difference between these properties and the laws of physics in non-virtual reality is that as far as we know, the laws of physics cannot be easily broken. In certain racing games, for example, specific glitches make it possible to skip vast amounts of the track and still 'complete' a lap according to the game (Summoning Salt 2018; 2022). Therefore, when constructing hard-coded laws and

regulations for our virtual society, preventing glitches in the system, and their exploitation, is essential. At the same time, it would be wise not to set boundaries that are too restrictive, or have too many procedures be hard-coded into the virtual society. If we did so, we would risk leaving future generations of citizens with an antiquated and inflexible set of restrictions, unable to be altered in light of future ideas and values.

With these two considerations in mind, let us consider what specific actions should be guarded against immutably. The example of a tyrant gaining control of a simulation to trap and torture everyone within it is a great starting point to find such actions. Naturally, we never want citizens to be trapped within the virtual society, just as we do not want citizens trapped in any particular country. Unless some temporary confinement measure is in place for a serious offense, a citizen should always be free to move between the virtual society and non-virtual reality as they please. Additionally, we do not want any citizen to experience pain unless they wish to incorporate it into a particular experience, such as a fantasy battle. Even in such cases, citizens should always retain control of how much pain or other unpleasantness they wish to experience at any given time.

In previous sections, I have described how the virtual society might be subdivided into different areas. These areas are the 'overworld'; where people interact and communities develop through collective decision-making, the private 'mini-universes' for each individual that can be modified in any way they choose, and the 'maturity worlds' for children to live in and grow naturally with their caretakers. The existence of these areas, their fundamental functions, and the ability of citizens to access them should also be made immutable facets. This is not to say that every aspect of those areas ought to be unchangeable by future governments. Rather, I believe that private spaces for all citizens, a wider communal space, and appropriate spaces for children to develop are all essential for the kind of virtual society we are looking to build. We can think of this as akin to the rights given to citizens of a nation through its constitution or bill of rights. In general, such elements of a nation are deliberately designed to be difficult to alter, such as the United States making amendments to its constitution an "onerous" process "in order to prevent arbitrary changes" (Obama White House Archives n.d.). In keeping with the spirit of a bill of rights, it is also important to make certain human rights immutable too, in order to prevent potential tyrants from disrupting free communication, doling out unwarranted punishments, or other authoritarian acts.

Now, let us explore what these immutable safeguards might mean in practice for our virtual society. Boese et. al. (2023, pg. 18) note two forms of "'democratic resilience' [defined as] the ability to prevent substantial regression in the quality of democratic institutions and practices". The first is 'onset resilience', where a state is able to prevent "declines in [its] democratic qualities", while 'breakdown resilience' refers to the state's ability to maintain its democracy in spite of such declines (Boese et. al. 2023, pg. 18). The study in question found that "Judicial constraints are positively and significantly associated with resilience to the onset of autocratization and to democratic breakdown once autocratization has begun", and that "Judicial institutions seem to play an important role as democracy's last line of defence against aspiring dictators" (Boese et. al. 2023, pg. 32). Other factors that were identified as positive contributors to democratic resilience include the presence of a parliamentary system as opposed to a presidential system, a "higher level of economic development", a long history of democracy, and a close proximity to other democracies (Boese et. al. 2023, pg. 32).

While our virtual society has no need for an internal economy, the high living standards it provides would resemble an exceptionally high level of economic development in non-virtual reality, and would likely produce the same boost to democratic resilience. Developing a strong judicial system is a more complicated matter, but the virtual society can aid in its implementation through the use of

specific rules, or ‘parameters’, that are incorporated into the programming. For example, if someone was required to attend a virtual trial, they could simply be teleported into the courtroom from wherever they may be. Some more pertinent examples for our purpose of preventing tyranny could include the inability to terminate a citizen’s employment without due process, or the inability to put legislation into place that falls substantially out of line of certain criteria for maintaining a balance of power.

Given that our virtual society encompasses both virtual and non-virtual space, a viable judiciary would need to do the same. This means that the virtual society will require at least a portion of its citizens to work primarily in non-virtual reality, not only for maintaining and defending the state against external threats, but internal ones as well. For example, if a malicious citizen (hereafter referred to as ‘Renholm’) wished to exit virtual reality and begin tampering with the hardware keeping it running, the virtual society would need guards and other security systems in non-virtual reality to detain him.

The dilemma that these external citizens cause is that they themselves could establish themselves as tyrants, given their position above and unconstrained by the virtual reality they are responsible for. While one Renholm acting alone in such a position may be stopped, an external group full of Renholms would be more able to operate without consequence, devastating the virtual society. To counteract this, a few measures can be implemented in tandem. These include:

1. Ensuring that those who work in non-virtual reality still live in virtual reality outside of their work. This ensures that those who maintain the virtual reality reap the benefits of their work, disincentivizing them from tampering and potentially making their own living situation worse. Additionally, this reduces the social separation between those who work in the virtual environment and those who work outside of it, increasing empathy between the two groups.
2. Develop a system whereby any changes made to the virtual society are communicated in easy-to-understand language to all citizens, and require the consent of a majority in order to come into effect.
3. Grant certain public servants who live entirely within the virtual society a certain degree of control over non-virtual reality. This control could include automated security systems that prevent potential Renholms working in non-virtual reality from disrupting hardware or breaking other laws.
4. Distributing power across all branches of government and necessary technological industries between numerous, diverse, and accountable stakeholders, as per the next AVAP that I will discuss.

### **AVAP 3 - Numerous Stakeholders and Governmental Structure**

One aspect of the virtual society we have yet to touch on is how it will be governed. There are many different ways a government can be structured, and choosing a model that is appropriate both for our virtual society and as a means to prevent tyranny is essential. I will begin by speculating on the kind of broader political landscape the virtual society may develop, then use that framework to create a governmental structure that allows a virtual utopia to remain politically flexible, without being politically unstable.

Firstly, it is worth addressing my reasons for preferring a politically flexible virtual utopia over a rigid and immovable one. As mentioned earlier, possible changes in the attitudes and desires of citizens over time should be taken into account. In addition, the non-virtual reality that the virtual society is

contingent upon will almost certainly continue changing, both diplomatically and circumstantially. Having a government that is able to adapt quickly to these changing external conditions is essential, as is ensuring that those adaptations best serve the wishes and wellbeing of the citizens. Lastly, a stagnant and unchanging society would likely be anathema to excitement and political engagement from its citizens. It would instead resemble the 'boring' utopias of classic literature that Finkielstein (2016, pg. 107) described; where common traits such as fixed schedules, standardisation, and inflexible community layouts "virtually guarantees psychologically debilitating sameness".

Now that my reasons are covered, let us examine the political landscape of our virtual society so far. Earlier, I noted that the virtual society would be made up of dozens of smaller communities within the same vast and malleable communal space. To support the flexibility and diversity of these communities, our virtual government may choose not to be responsible for how each community governs itself, only intervening if certain rights are violated. For example, a community that wishes to style itself as a medieval kingdom may be free to conduct itself in this manner, as long as there are no violations of consent or attempts to prevent people from leaving the community.

Therein lies the first political axis for which citizens of the virtual society might find themselves on; the level of regulation for virtual communities. Some citizens could find it repugnant for the government to allow certain practices even in a consensual and non-lethal manner, such as simulated trial by combat in a medieval-themed community. For others, it might be a less ideological matter. If someone lived in a neon noir-themed community, they might simply prefer not to see or hear a jovial candyland community nearby, much like how a tree on one person's property might cause consternation due to it obstructing the view from another person's balcony. And of course, some citizens may support lower levels of regulation, just as Oxford locals in 1986 supported the continuation of an unconventional shark sculpture jutting out of a rooftop, despite the city council citing safety concerns and a lack of planning permission (Purves 2007).

Another political subject in the virtual society would undoubtedly be the non-virtual reality. As noted earlier, there are a plethora of reasons as to why virtual citizens would be invested in non-virtual affairs. In the political realm, these could include diplomacy, resource extraction, and the maintenance of crucial non-virtual systems. The last issue that I believe will constitute a major political axis in the virtual society is mental wellbeing. As noted in Section 2, living your life in a utopia, unburdened by the need to survive, increases the importance of you actively seeking out purpose, fulfillment, and meaningful relationships. Therefore, it matters how the virtual government aids individuals in those pursuits, and some citizens will likely have strong feelings about that. We can imagine the discourse might be similar to how people in modern society debate social welfare programs and other measures like it (Blendon et. al. 1995; Lundberga & Syltevik 2024), albeit with a greater focus on how much the government ought to involve itself in meaning-making, and in what manner.

We could expand our range of possible political issues further if we wish, including debates about the contents of maturity worlds or the process of acquiring or revoking virtual citizenship, but the point is clear. Even in a virtual utopia, citizens would undoubtedly form strong opinions about, and engage with, the political system. To deny individuals the ability to meaningfully engage may foster a sense of powerlessness at best, and at worst may encourage them to take forceful action. If we wish to avoid instability and revolution, our virtual government needs a mechanism for citizens to meaningfully interact with it, have their thoughts be represented, and hold the government to account when justified.

With this in mind, I propose that our virtual government ought to incorporate both direct democracy and sortition. Sortition and direct democracy are defined by Christiano and Bajaj (2024, n.p.) as "the

appointment of political officials by random selection”, and “direct referenda of the members of a society in deciding on the laws and policies of the society”, respectively.

My first reason for choosing such a combination is that it has been utilized before, most notably in the classical city-state of Athens (Christiano & Bajaj 2024). In addition, based upon personal experience, many virtual communities across the modern internet already look towards their users for feedback and policy decisions, a process akin to direct democracy. For example, one subreddit (an online group managed by reddit users via the Reddit website) was being inundated with AI-generated artwork, and the moderators of that subreddit allowed its users to vote on whether or not such artwork should continue to be allowed on the subreddit. Following the vote, AI-generated art was banned, in line with the community’s wishes. Even in far larger communities, such as the Minecraft player base, the developers of the game incorporate a significant amount of community feedback into their development process (Cooper 2024).

My second reason for favoring direct and sortition-based democracy over the more common representative democratic system is that certain aspects of representative democracy may serve to incentivise abuses of power in the virtual society. For one, there is evidence that the subsection of individuals who put themselves in the running for elected leadership positions, particularly on a national level, have a greater proportion of traits we are aiming to avoid; namely narcissism and psychopathy, compared to the general population (Chen et. al. 2021). A government derived from sortition, where the sample of selected individuals matches the wider population, would naturally have reduced levels of such traits by comparison. Of course, there are several ways that ‘matching’ with the wider population could be carried out poorly, and it would be important to ensure that the process includes valuable minority perspectives, and can keep oppressive ideologies at arms length.

Furthermore, it is easier for malicious actors, both inside and outside of a society, to influence politics and policy when they know which organizations and people to influence. Political parties are very common in representative democracies, and malicious actors who can aid in the electoral success of those parties, for example through donations or promotional materials, could extract favors or policy concessions from them. By comparison, a virtual society governed through sortition has no need for political parties, nor is there anything for malicious actors to offer individuals in government. Promotion and popularity have no bearing on who is selected to serve in government, and bribery is essentially useless, since everyone already has access to any experience they desire.

My third reason for encouraging sortition and direct democracy is that it diversifies and distributes power across as many individuals as possible, aiding further in the prevention of power abuse. Thanks to direct democracy, the entire population will have a say in all major legislation, while a government selected by ongoing sortition ensures that the political class is dynamic and less entrenched in any particular kind of individual. Additionally, those with the technological and scientific expertise required to keep the virtual society running will also serve in an advisory role for the judicial and legislative branches, and have a portion of executive power similar to that of federal civil servants in the United States (Government Organization and Employees 1966). This ensures that such individuals serve the virtual society in accordance with their abilities, and provide an additional check on abuses of power, without being able to easily abuse their own power.

My final reason is that a government rooted in sortition and direct democracy helps fulfill the earlier AVAP of disincentivization. In that section, I described the importance of careers being viewed as responsibilities and a means to serve your community, instead of opportunities for oneself. A great example of this in our current age is jury duty, where citizens are selected to act as jurors in a court case and ultimately deliver a verdict. Naturally, we cannot randomly select individuals for every position. Even in a virtual society, numerous vital positions would require experts to fill them.

However, it is reasonable to assume that a random selection process for the legislative and executive, as is done to select a jury, would foster a perception of those roles similar to that of jury duty; a responsibility.

The component of direct democracy also plays a role in disincentivization, particularly when it comes to revolution. As I noted near the start of this section, it is essential that all citizens have a means to engage with the political process, even in a virtual utopia. Otherwise, we risk citizens feeling politically powerless, increasing the likelihood of dangerous measures being used to regain political power. Direct democracy grants a significant means of political engagement, and multiple studies have found it capable of boosting both political participation and trust in political institutions (Christensen 2019; Tolbert et. al. 2003). Therefore, even though citizens will not vote for their leaders, they will still be able to play an important role in the virtual society's democratic institutions.

## **AVAP 4 - Transparency and Privacy**

A hallmark of many dystopian futures is the constant and insidious monitoring of citizen's lives, with *Nineteen Eighty-Four* and *Brave New World* being notable works that feature an extreme disregard for privacy (Hinchliffe 2021). The concern of malicious actors gathering data on citizens is particularly salient in the case of a virtual society. Already in our modern era, it has been known for several years that various technology companies have been tracking and selling the browsing data of their users to advertisers (Olejnik et. at. 2013; Wu 2024). In a virtual society, the very brain activity of its citizens may be constantly monitored by algorithms and artificial intelligences in order for the simulation to work. Even if such activity is not being monitored maliciously or sold to unknown parties, this level of invasiveness is understandably alarming.

Fortunately, the complexity of such a system may prove to be a benefit for privacy. Many artificial intelligences in modern times, including ChatGPT and Alexa, can be described as 'Black Box AI', meaning that their "internal workings and decision-making processes are not transparent, making it unclear how they arrived at their conclusions", even for the developers of the software (Glover 2024, n.p.). For example, you could ask ChatGPT to describe an imaginary cabin in the woods to you, and through an algorithmic process unknown to all, it could do so. Even if someone were to peek into the code without knowing either your input ("Describe a hypothetical cabin in the woods") or the output (the answer ChatGPT gives you), it would be very difficult if not impossible for them to work out what those inputs and outputs were.

An artificial intelligence capable of reading and understanding someone's brain activity, and using that information to recreate a full sensory experience of a consistent virtual world, all in real time, would be orders of magnitude more advanced and complex than even our most powerful artificial intelligences today. Given that we can already create complex algorithms that are opaque to even their developers, a much more complex system could be made even more opaque, rendering its code impossible for malicious humans to extract personal information from. It is important to note that, unlike the current algorithms that provide meaningful user data to advertisers, the virtual reality programs would not be designed with the intent to transmit their outputs to a third party. In an ideal virtual society, the output would only be experienced within the mind of the individual.

Of course, this black box aspect of certain AI is not without its issues. Warren J. von Eschenbach (2021, pg. 1608) describes the well-known "black box problem in AI", where the lack of transparency and subsequent unpredictability of such systems erodes our capacity to trust them, a trust that is important to establish in a world increasingly reliant upon these particular sorts of AI. This concern

would be even more salient in our virtual society. However as Eschenbach (2021, pg. 1620) points out, 'black box AI' do not exist independently of the broader context they operate under, both in terms of the systems that work in tandem with them, and the "larger context of institutional or organizational norms and standards that safeguard the interests and goods of those [they serve]".

For an example, let us think of a search engine like Google, using a black box AI to provide us the best results possible from our searches. While it may be the case that someone does not trust Google's AI exclusively, it seems more likely that they also do not trust Google as a company to develop a black box AI that ensures the best experience for its users. They may believe Google would cut corners during development, making the AI sloppy, or design the AI to prioritize engagement and profit over accuracy of information. As a result, one may be more likely to trust Google's black box AI's if Google itself is able to earn that trust. More broadly, "companies and institutions that develop and use [AI] and who have great control over the socio-technical system need to ensure that they themselves earn our trust" (Eschenbach 2021, pg. 1620).

This would undoubtedly be one of the methods a virtual society takes to ensure that its citizens can trust the AI it relies on, perhaps in conjunction with yet unknown methods of building safer black box AI. Assuming this can be done, the opacity and security of this system could be made an immutable facet of the virtual society. Not only would each person's experiences be inaccessible to all others from a coding perspective, but certain privacy-violating simulations could be blocked as well. For example, if you and Renholm both lived in a virtual society, and Renholm requested that the interior of his house "perfectly resemble" your own, the program might deny the request on privacy grounds. Repeated attempts with slightly different wording, like "almost perfectly" or "with the following miniscule changes" would still be denied, and could even result in Renholm being flagged for attempted privacy violations.

The mechanism for blocking certain types of content is possible even in the large language models of today. If I ask ChatGPT how to make a bomb, it will refuse to answer. The same is true if I ask it to write a transcript for an educational video on the subject. Even if I ask it to write a story about a man who builds a bomb, the story will not contain any useful details about the process itself. This has remained true across multiple tests I have done on multiple ChatGPT models throughout 2024 and 2025. These safeguards are vital for ensuring that harmful and private information is not easily accessible, just as they would be in a virtual society.

However, there are some pieces of information that are in the public interest, even if they may be detrimental to someone's reputation. In particular, for this section on abuses of power, I am referring to the need for governmental and legal transparency. Legal transparency, or more accurately legal accessibility, is the principle that law, among other things, "should be a body of norms promulgated as public knowledge so that people can study it, internalize it, figure out what it requires of them, and use it as a framework for their plans and expectations and for settling their disputes with others" (Waldron 2023, n.p.). This precedent helps ensure that citizens are less likely to be blind-sided by themselves or others breaching laws that they were unaware of, decreasing the likelihood of citizens being distrustful or resentful towards their government or judicial system.

Transparency, in regards to government, is defined by Johnston (2024, n.p.) as "[the] capacity of outsiders to obtain valid and timely information about the activities of government". An example of this is Parliament on Demand, a service in New Zealand that provides live footage of its parliamentary sessions (Parliament on Demand n.d.). There are a myriad of reasons to increase government transparency in modern societies, and those reasons are largely consistent with a virtual society too. First of all, it allows citizens to better hold the government to account, and speak truth to power. For example, if citizens were aware that new legislation was being discussed that

would allow virtual communities to roleplay as fascist states, they would have a longer amount of time to communicate their thoughts and engage with others about the issue before it came to a vote.

Political commentary, parody, and satire have already been made much easier to create and distribute for both ordinary people and television networks thanks to the rise of the internet, video sharing, and social media. While it is true that misinformation may also spread further as a result, those interested in genuinely informing, satirizing, or advocating have also seen tremendous success doing so online, including John Oliver (LastWeekTonight n.d.), Harry Brewis (hbombguy n.d.), the Onion (Purcell 2017), Leeja Miller (n.d.), the ASAN (Autistic Self Advocacy Network n.d.), and various news organizations such as TLDR News (n.d.) and the Associated Press (n.d.), amongst many others. In a virtual society where everyone has the ability to create and showcase any experience, the potential for genuine advocacy and informational content would be even greater. Government transparency can aid in both supporting that kind of necessary content, and in fostering a level of trust between the government and the wider population.

## **AVAP 5 - Emergency Measures**

Lastly, it would be prudent to briefly consider backup plans in the event that virtual abuse is somehow able to run rampant in spite of the four other AVAPs. For the virtual society, I would propose a two-tiered approach to such a system, each tier composed of multiple possible methods of countering abuses of power. The first tier would serve to bolster the previously mentioned 'breakdown resilience' (Boese et. al. 2023), alongside the society's judicial system and strong democratic foundation. The second tier would serve as a 'last resort' system, one that prioritizes the evacuation or safety of citizens over maintaining the society itself. It is worth noting that I will be presenting the methods in each tier as possibilities rather than definitive strategies. In addition, it is important to be careful when developing these or any other emergency measures one may think up, to ensure that they themselves do not become a means of spreading tyranny.

### Tier 1 Proposals

*Citizen Impeachment* - The immutable capacity of citizens outside of governmental roles to revoke the position of representative from one or more individuals. This would be on the condition that those citizens, as a group, can demonstrate flagrant disregard for key principles of the virtual society on the part of the representative(s), and that those citizens are a significant portion of the total population.

*Countering Misinformation* - This counter utilizes the virtual nature of the society to identify and correct misattribution. The system scans citizen-created media for anything that is both presenting as non-fictional and claiming to be from a source that it is not actually from. For example, a fake speech from a political figure, disseminated as if it actually took place, would count in this scenario. The system adds a marker to each instance of that media to let citizens watching it know that the media content has been misattributed. In addition, the original creator of the content could be flagged for spreading misinformation. These counters would only apply to content that is spread publicly. Private works would be unaffected.

### Tier 2 Proposals

*Evacuation* - This procedure comes into effect if the system, either through an internal or external evaluation mechanism that routinely scans for certain issues, detects that citizens can no longer exit the virtual society of their own accord. If that is the case, the system triggers an evacuation,

disconnecting citizens from the virtual society and granting them the means to relocate of their own volition.

*Reboot* - If the same system detects a prolonged period of suffering across the population, or if an evacuation command fails, the system could be forced to reboot to its original settings, clearing any changes made or leadership positions currently filled. If a reboot fails, then a targeted shutdown of everything except the life-support systems could come into effect, keeping citizens safe until help arrives from the outside world.

## **Summary of Section 3**

The factors of a virtual society that contribute positively to wellbeing are meaningless if those factors can be diminished or even dismantled entirely. As such, it is vital that a conversation on how virtual societies can be beneficial addresses the ways a virtual society can remain beneficial as well. This section has addressed how virtual societies can be very useful in disincentivizing abuses of power, as well as how they can provide legal frameworks even more robust than our current ones. Additionally, the creation of a virtual society from the ground up grants us the opportunity to reevaluate the ways in which democracies are run, in order to improve engagement and lessen corruption.

## **Conclusion**

### **Summary**

Building upon the groundwork laid out in Chalmers' Reality+ (2022), we have addressed a number of prudential concerns regarding the impact a futuristic virtual life could have on wellbeing, not only on an individual level, but on a societal level as well. The valuable aspects of non-virtual reality that we fear are missing from virtual reality may actually be present when a virtual society is designed carefully. In particular, the concerns of a virtual society being unfulfilling, monotonous, or hindering natural human development, may not apply to well-designed virtual societies. In fact, virtual reality may be better equipped to tackle said issues than non-virtual reality is.

In addition, we have gradually devised an internally consistent vision for our hypothetical virtual society, one that satisfies our criteria for contributing positively to individual wellbeing. To ensure that this society would not fall victim to abuses of power, five preventative axioms were able to be incorporated into this vision too. These axioms bolster the virtual society's resilience without impinging upon citizens personal liberty and their rights to privacy and political participation. The vision as a whole provides an initial model society to examine, pick apart, and refine further.

Overall, this exercise has given us reason to rethink the idea that a fully immersive virtual life would be intrinsically and automatically worse for someone's well being than a life in non-virtual reality.

### **Potential Implications**

Humanity is a long way from being able to construct a fully-immersive virtual society. However, by exploring the potential of the concept and highlighting it in a positive manner, it may encourage further dialogue about the benefits of virtual reality. Given the pervasive sentiment of a virtual life as dystopian, this alternate perspective will become more and more valuable as technology ever

improves. More importantly, an equal public sentiment between positive and negative associations may be more conducive to finding a middle ground for the technology's usage, rather than an outright disregard for it that a largely negative sentiment would likely entail.

Much of this is speculative of course, and assuming that this research will shift perception in any way, which is not a guarantee. Something to stress is that future discussions ought to keep in mind the concerns of virtual reality and find means to address them, rather than embracing any form of virtual society that may present itself. This is an important aspect of the philosophical discussion about virtual reality; ideally research like this can steer that conversation in a balanced direction, where the potential of the technology is not dismissed outright, without losing sight of the objective of improving wellbeing.

## **Areas of Further Inquiry**

Given that the virtual society is only one possible future out of many, it is worth examining within that broader context. We might find that compared to other potential futures, the virtual society could be less valuable to wellbeing, or more. It might be more feasible in some ways and less so in others. It could be an end goal for humanity, or a stepping stone to a better future. Another question would be the steps necessary to move society towards such a project, if it turns out to be in our best interest. Not only would we require far greater virtual reality technology than we currently possess, but we would likely need to make considerable changes to our political, economic, and value systems first. The wrong conditions by which the virtual society arises could have disastrous consequences for its citizens.

Developing any sort of complex society from the ground up naturally requires a great deal of care and attention to detail, and this is particularly true when designing a virtual society. The technological and experiential aspects are key components, without which the society does not work. If a resilient virtual utopia is to be realised, both the overarching design principles and the minutiae of such a project must be discussed and examined thoroughly, across multiple disciplines ranging from philosophy and political science to biology, computer science, psychology, and even fields such as game design.

There is also the technology and virtual societies of the here and now to consider as well. I have mentioned both VRChat and Decentraland as examples of modern communities in virtual reality. It may be the case that the first virtual society, rather than being developed from scratch in the far future, will gradually form out of a current or near-future virtual community such as theirs. If that is the case, then the examination of such communities and their potential development is essential, especially in regards to the well-being and sense of community that individuals involved experience. If a cultural shift appears to be manifesting in such communities, this may be pivotal in determining what future virtual societies could look like.

## References

- Asselin, C. (2023, May 24). Social Media Giants Fight Against Fake News: Moving Forward or Stagnant? *Digimind*. <https://blog.digimind.com/en/insight-driven-marketing/social-media-giants-fight-fake-news>
- Associated Press. (n.d.). *About Us*. <https://www.ap.org/about/>
- Autistic Self Advocacy Network. (n.d.). *About ASAN*. <https://autisticadvocacy.org/about-asan/>
- Auxier, B., Anderson, M., Perrin, A., & Turner, E. (2020, July 28). Parenting children in the age of screens. *Pew Research Center*. <https://www.pewresearch.org/internet/2020/07/28/parenting-children-in-the-age-of-screens/>
- Auxler, B., & Arbanas, J. (2023, July 27). While we wait for the metaverse to materialize, young people are already there. *Deloitte Insights*. <https://www2.deloitte.com/us/en/insights/industry/technology/gen-z-and-millennials-are-metaverse-early-adopters.html>
- Blastbolt the Bolt. (2023, April 27). *Getting Over It Speedrun World Record in 59.885s* [Video]. YouTube. <https://www.youtube.com/watch?v=cb8yNvln8Q>
- Blendon, R. J., Altman, D. E., Benson, J., Brodie, M., James, M., & Chervinsky, G. (1995). The public and the welfare reform debate. *Archives of pediatrics & adolescent medicine*, 149(10), 1065-1069. <https://doi.org/10.1001/archpedi.1995.02170230019002>
- Boese, V. A., Edgell, A. B., Hellmeier, S., Maerz, S. F., & Lindberg, S. I. (2023). How democracies prevail: democratic resilience as a two-stage process. In *Resilience of Democracy* (pp. 17-39). Routledge. <https://doi.org/10.1080/13510347.2021.1891413>
- Brickman, P., Coates, D., & Janoff-Bulman, R. (1978). Lottery winners and accident victims: Is happiness relative?. *Journal of personality and social psychology*, 36(8), 917. <https://doi.org/10.1037/0022-3514.36.8.917>
- Carlyle, T. (1843, July 31). Past and Present - Thomas Carlyle's Collected Works, Vol. XIII. *Gutenberg*. <https://www.gutenberg.org/files/26159/26159-h/26159-h.htm>
- CDC. (2024, May 15). *Health Effects of Social Isolation and Loneliness*. <https://www.cdc.gov/social-connectedness/risk-factors/index.html>
- CGP Grey. (2016, October 25). *The Rules for Rulers* [Video]. YouTube. <https://www.youtube.com/watch?v=rStL7niR7gs>
- Chalmers, D. J. (2022). *Reality+: Virtual worlds and the problems of philosophy*. Penguin UK.
- Chen, P., Pruyers, S., & Blais, J. (2020). The Dark Side of Politics: Participation and the Dark Triad. *Political Studies*, 69(3), 577-601. <https://doi.org/10.1177/0032321720911566> (Original work published 2021)
- Child Health Evaluation and Research Centre. (2023, October 16). *New from the Mott Poll: Many parents of elementary aged children may engage in more helicopter parenting than they think*. <https://chea.org/news-publications/2023-10-16/mott-poll-independence>

- Chin, M. (2023, September 13). Apple's Vision Pro headset is on track to ship early next year. *The Verge*. <https://www.theverge.com/2023/9/12/23860400/apple-vision-pro-mixed-reality-headset-release-date>
- Christensen, H. S. (2019). Boosting political trust with direct democracy? The case of the Finnish citizens' initiative. *Politics and Governance*, 7(2), 173-186. <https://doi.org/10.17645/pag.v7i2.1811>
- Christiano, T., & Bajaj, S. (2024, June 18). Democracy. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/democracy/>
- Cooper, R. (2024, September 8). The future of Minecraft's development. *Mojang*. <https://www.minecraft.net/en-us/article/the-future-of-minecrafts-development>
- Curtis, G. (2020). Appeal to nature. *Fallacy Files*. <https://effectiviology.com/appeal-to-nature-fallacy>
- Davenport, C. (2023, August 19). As Russia and India attempt lunar landings, the moon rush gets crowded. *Washington Post*. <https://www.washingtonpost.com/technology/2023/08/19/russia-india-moon-landing-artemis/>
- De Brigard, F. (2010). If you like it, does it matter if it's real?. *Philosophical Psychology*, 23(1), 43-57. <https://doi.org/10.1080/09515080903532290>
- Deighan, M. T., Ayobi, A., & O'Kane, A. A. (2023, April). Social virtual reality as a mental health tool: how people use VRChat to support social connectedness and wellbeing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-13).
- Ditzler, G. (2018, February 17). Getting Over It with Bennett Foddy Review: Why Must You Hurt Me. *GameSkinny*. <https://www.gameskinny.com/reviews/getting-over-it-with-bennett-foddy-review-why-must-you-hurt-me/>
- Ettman, C. K., Adam, G. P., Clark, M. A., Wilson, I. B., Vivier, P. M., & Galea, S. (2022). Wealth and depression: A scoping review. *Brain and Behavior*, 12(3), e2486. <https://doi.org/10.1002/brb3.2486>
- Finkielsztein, M. (2016). Boredom and melancholy in utopias and dystopias. In *More After More. Essays Commemorating the Five Hundredth Anniversary of Thomas More's Utopia* (pp. 99-112). Ośrodek Badawczy Facta Ficta.
- Folding Ideas. (2023, March 27). *The Future is a Dead Mall - Decentraland and the Metaverse* [Video]. YouTube. <https://www.youtube.com/watch?v=EiZhdpLXZ8Q>
- Forcehimes, A. T., & Semrau, L. (2016). Well-Being: Reality's Role. *Journal of the American Philosophical Association*, 2(3), 456-468. <https://doi.org/10.1017/apa.2016.23>
- Forte, T. (2018, April 11). Treconomics: The Economics of Post-Scarcity. *Forte Labs*. <https://fortelabs.com/blog/treconomics-productivity-in-a-post-scarcity-economy/>
- Fouk, T. A., Chighizola, N., & Chen, G. (2020). Power corrupts (or does it?): An examination of the boundary conditions of the antisocial effects of experienced power. *Social and Personality Psychology Compass*, 14(4), e12524. <https://doi.org/10.1111/spc3.12524>
- Froebel, F. (1886). *The education of man* (Vol. 5). A. Lovell & Company.
- Glover, E. (2024, August 6). What Is Black Box AI? *Built In*. <https://builtin.com/articles/black-box-ai>

- Government Organization and Employees, 5 U.S.C. § 1101 – 1508 (n.d.).  
<https://www.law.cornell.edu/uscode/text/5>
- Grand View Research. (2023). *Virtual Reality (VR) Market Size, Share & Trends Analysis Report By Technology (Semi & Fully Immersive, Non-immersive), By Device (HMD, GTD), By Component (Hardware, Software), By Application, By Region, And Segment Forecasts, 2023 – 2030* (Grand View Research market analysis report #GVR-1-68038-831-2)
- Gregston, M. (2020, February 28). Fear of Failure Vs. Freedom to Learn. *Parenting Today's Teens*.  
<https://parentingtodaysteens.org/articles/fear-of-failure-vs-freedom-to-learn/>
- Gruebner O., Rapp M. A., Adli M., Kluge U., Galea S., & Heinz A. (2017). Cities and mental health. *Deutsches Ärzteblatt International*, 114(8), 121.  
<https://doi.org/10.3238/arztebl.2017.0121>
- Guardian staff and agency. (2022, July 23). Google fires software engineer who claims AI chatbot is sentient. *The Guardian*. <https://www.theguardian.com/technology/2022/jul/23/google-fires-software-engineer-who-claims-ai-chatbot-is-sentient>
- Gunnell. (2022, November 7). An Elon Musk chatbot tells Insider he wants to buy CNN, reinstate Trump on Twitter, and 'show people how the sausage gets made'. *Business Insider*.  
<https://www.businessinsider.com/artificial-intelligence-elon-musk-twitter-epstein-biden-ye-chatbot-2022-11>
- Guterman, D. (Executive Producer). (2023). *Carol & the End of the World* [TV series]. Netflix Animation; Netflix.
- hbombberguy. (n.d.). *Home* [YouTube channel]. YouTube. Retrieved February 5, 2025, from <https://www.youtube.com/@hbombberguy>
- Haidt, J. (2024). *The anxious generation: How the great rewiring of childhood is causing an epidemic of mental illness*. Penguin.
- Halbrook, Y. J., O'Donnell, A. T., & Msetfi, R. M. (2019). When and how video games can be good: A review of the positive effects of video games on well-being. *Perspectives on Psychological Science*, 14(6), <https://doi.org/10.1177/174569161986380>
- Hinchliffe, J. (2021). Speculative fiction, sociology, and surveillance studies: towards a methodology of the surveillance imaginary. *Surveillance & Society*, 19(4), 414-424.  
<https://doi.org/10.24908/ss.v19i4.15039>
- Hindriks, F., & Douven, I. (2018). Nozick's experience machine: An empirical study. *Philosophical Psychology*, 31(2), 278-298. <https://doi.org/10.1080/09515089.2017.1406600>
- Hunter, N. (2025, March 18). Every Meta Quest Model: A Full History of Release Dates. *IGN*.  
<https://www.ign.com/articles/all-meta-quest-models-in-order>
- Hurley, M. (Writer), & Conway, J. L. (Director). (1988, May 16). The Neutral Zone (Season 1, Episode 26) [Television series episode]. In G. Roddenberry (Executive producer), *Star Trek: The Next Generation*. Paramount; various
- Jackson, F. (1986). What Mary didn't know. *The journal of philosophy*, 83(5), 291-295.  
<https://doi.org/10.2307/2026143>

- Johannes, N., Vuorre, M., & Przybylski, A. K. (2021). Video game play is positively correlated with well-being. *Royal Society open science*, 8(2), <https://doi.org/10.1098/rsos.202049>
- Johnston, M. (n.d.). Transparency. *Britannica*. <https://www.britannica.com/topic/transparency-government>
- JustWatch. (n.d.). *All Movies and TV Series on Netflix*. <https://www.justwatch.com/uk/provider/netflix>
- Kato, T. A., Kanba, S., & Teo, A. R. (2019). Hikikomori: Multidimensional understanding, assessment, and future international perspectives. *Psychiatry and clinical neurosciences*, 73(8), 427-440. <https://doi.org/10.1111/pcn.12895>
- Kawachi, I., Adler, N. E., & Dow, W. H. (2010). Money, schooling, and health: Mechanisms and causal evidence. *Annals of the New York Academy of Sciences*, 1186(1), 56-68. <https://doi.org/10.1111/j.1749-6632.2009.05340.x>
- Kolber, A. J. (1994). Mental Statism and the Experience Machine. *Bard Journal of Social Sciences* 3, 10-17.
- Lassiter, C., & Kagan, A. (2024). *Review of David Chalmers, Reality+: virtual Worlds and the problems of Philosophy*, New York: WW Norton & Company, 2022.
- LastWeekTonight. (n.d.). *Home* [YouTube channel]. YouTube. Retrieved February 5, 2025, from <https://www.youtube.com/@LastWeekTonight>
- Leeja Miller. (n.d.). *Home* [YouTube channel]. YouTube. Retrieved February 5, 2025, from <https://www.youtube.com/@LeejaMiller>
- Lindqvist, E., Östling, R., & Cesarini, D. (2020). Long-run effects of lottery wealth on psychological well-being. *The Review of Economic Studies*, 87(6), 2703-2726. <https://doi.org/10.1093/restud/rdaa006>
- Lundberg, K. G., & Syltevik, L. J. (2024). Welfare stigma in a social democratic welfare regime during a decade of national public debate: production, contestation and continuities. *European Societies*, 26(5), 1572-1595. <https://doi.org/10.1080/14616696.2024.2333526>
- McClean Hospital. (2024, December 27). *Do You Manage Your Time Well?* <https://www.mcleanhospital.org/essential/do-you-manage-your-time-well>
- McMillan, K. A., Enns, M. W., Asmundson, G. J., & Sareen, J. (2010). The association between income and distress, mental disorders, and suicidal ideation and attempts: findings from the Collaborative Psychiatric Epidemiology Surveys. *The Journal of clinical psychiatry*, 71(9), 1497. <https://doi.org/10.4088/JCP.08m04986gry>
- Montemorano, C. (2020). *Body Language: Avatars, Identity Formation, and Communicative Interaction in VRChat* [Honours Project]. Eagle Scholar. [https://scholar.umw.edu/student\\_research/361/](https://scholar.umw.edu/student_research/361/)
- Mukherjee, U. (2021). Navigating children's screen-time at home: Narratives of childing and parenting within the familial generational structure. *Children's Geographies*, 19(6), 646-658. <https://doi.org/10.1080/14733285.2020.1862758>

- Nai, A., & Toros, E. (2020). The peculiar personality of strongmen: Comparing the Big Five and Dark Triad traits of autocrats and non-autocrats. *Political Research Exchange*, 2(1), 1707697. <https://doi.org/10.1080/2474736X.2019.1707697>
- Nozick, R. (1974). *Anarchy, state, and utopia*. New York: Basic books.
- Obama White House Archives. (n.d.). *The Constitution*. <https://obamawhitehouse.archives.gov/1600/constitution>
- Occhipinti, J. A., Hynes, W., Geli, P., Eyre, H. A., Song, Y., Prodan, A., & Hickie, I. B. (2023). Building systemic resilience, productivity and well-being: a Mental Wealth perspective. *BMJ Global Health*, 8(9). <https://doi.org/10.1136/bmjgh-2023-012942>
- Oh, H. J., Kim, J., Chang, J. J. C., Park, N., & Lee, S. (2023). Social benefits of living in the metaverse: The relationships among social presence, supportive interaction, social self-efficacy, and feelings of loneliness. *Computers in Human Behavior*, 139, 107498-. <https://doi.org/10.1016/j.chb.2022.107498>
- Oh, V. Y. S. (2025). Money matters for future well-being: A latent growth analysis and meta-analytic integration of associations between income, financial satisfaction, and 22 well-being variables across three data sets. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspp0000552>
- Olejnik, L., Minh-Dung, T., & Castelluccia, C. (2013). Selling off privacy at auction. *HAL Open Science*, <https://inria.hal.science/hal-00915249/>
- Parliament on Demand. (n.d.). *About Parliament Video*. <https://videos.parliament.nz/about>
- Pascal, B. (1887). *The thoughts, letters, and opuscles of Blaise Pascal*. Houghton, Mifflin.
- Patel, V., Burns, J. K., Dhingra, M., Tarver, L., Kohrt, B. A., & Lund, C. (2018). Income inequality and depression: a systematic review and meta-analysis of the association and a scoping review of mechanisms. *World Psychiatry*, 17(1), 76–89. <https://doi.org/10.1002/wps.20492>
- Pierce, D. (2023, September 28). Meta's \$499.99 Quest 3 headset is all about mixed reality and video games. *The Verge*. <https://www.theverge.com/2023/9/27/23889059/meta-quest-3-headset-mixed-reality-price-release-date>
- Platts, K. (2020, March 24). Keeping Up: An Introduction to Speedrunning. *Redbrick*. <https://www.redbrick.me/intro-to-speedrunning/>
- Prins, S. J., Bates, L. M., Keyes, K. M., & Muntaner, C. (2015). Anxious? Depressed? You might be suffering from capitalism: contradictory class locations and the prevalence of depression and anxiety in the USA. *Sociology of Health & Illness*, 37(8), 1352–1372. <https://doi.org/10.1111/1467-9566.12315>
- Purcell, C. (2017, August 23). The Onion: inside the biggest 'fake news' outlet in the world. *The Sydney Morning Herald*. <https://www.smh.com.au/entertainment/the-onion-inside-the-biggest-fake-news-outlet-in-the-world-20170822-gy1chg.html>
- Purohit, A. K., Barclay, L., & Holzer, A. (2020, April 25). Designing for digital detox: Making social media less addictive with digital nudges. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-9). <https://doi.org/10.1145/3334480.3382810>

- Purves, L. (2007, August 9). Let's salute fibreglass fish and wacky artists. *The Times*.  
<https://www.thetimes.com/travel/destinations/uk-travel/england/london-travel/lets-salute-fibreglass-fish-and-wacky-artists-s2lp35j5tvb?region=global>
- Rummel, R. J. (2018). *Death by Government: Genocide and Mass Murder Since 1900*. Taylor and Francis. <https://doi.org/10.4324/9780203793756>
- Saura, J. R., Palacios-Marqués, D., & Iturricha-Fernández, A. (2021). Ethical design in social media: Assessing the main performance measurements of user online behavior modification. *Journal of Business Research*, 129, 271-281.  
<https://doi.org/10.1016/j.jbusres.2021.03.001>
- Schur, M. (Executive Producer). (2016-2020). *The Good Place* [TV series]. Fremulon; NBC.
- Schwartz, B. (2015). The paradox of choice. *Positive psychology in practice: Promoting human flourishing in work, health, education, and everyday life*, 121-138.  
<https://doi.org/10.1002/9781118996874.ch8>
- Speedrun.com. (n.d.). *Getting Over It With Bennett Foddy*. <https://www.speedrun.com/goiwbfb/stats>
- Stanton, A. (Director). (2008). *WALL-E* [Film]. Pixar Animation Studios; Walt Disney Pictures.
- Steam. (n.d.). *Global Gameplay Stats - Getting Over It with Bennett Foddy*.  
<https://steamcommunity.com/stats/240720/achievements>
- Steamcharts. (n.d.). *VRChat*. <https://steamcharts.com/app/438100>
- Stewart, F. (2014). Against happiness: A critical appraisal of the use of measures of happiness for evaluating progress in development. *Journal of Human Development and Capabilities*, 15(4), 293-307. <https://doi.org/10.1080/19452829.2014.903234>
- Straszfilms. (2021, February 4). *Identity, Gender, and VRChat (Why is everyone in VR an anime girl?)* [Video]. YouTube. [https://www.youtube.com/watch?v=5v\\_DI7i4Bcw](https://www.youtube.com/watch?v=5v_DI7i4Bcw)
- Summoning Salt (2018, August 27). *Mario Kart Wii: The History of the Ultra Shortcut* [Video]. YouTube. [https://www.youtube.com/watch?v=mmJ\\_LT8bUj0](https://www.youtube.com/watch?v=mmJ_LT8bUj0)
- Summoning Salt (2022, July 24). *The Hidden Shortcuts of Mario Kart Double Dash* [Video]. YouTube. <https://www.youtube.com/watch?v=ToO7leVx1y8>
- Tägil, S. (1998). *Alfred Nobel's thoughts about war and peace*. Nobelprize.org.
- Tale Foundry. (2024, January 19). *What If The World Never Ends?* [Video]. YouTube. <https://www.youtube.com/watch?v=ugdB-0mnE5A>
- Tan, J. J. X., Kraus, M. W., Carpenter, N. C., & Adler, N. E. (2020). The Association Between Objective and Subjective Socioeconomic Status and Subjective Well-Being: A Meta-Analytic Review. *Psychological Bulletin*, 146(11), 970–1020. <https://doi.org/10.1037/bul0000258>
- The Wildlife Trusts. (2021, December 1). *New poll reveals people think children need more time learning in nature*. <https://www.wildlifetrusts.org/news/poll-children-need-time-nature>
- TLDR News. (n.d.). *Home* [YouTube channel]. YouTube. Retrieved February 5, 2025, from <https://www.youtube.com/@TLDRnews>

- Tolbert, C. J., McNeal, R. S., & Smith, D. A. (2003). Enhancing civic engagement: The effect of direct democracy on political participation and knowledge. *State Politics & Policy Quarterly*, 3(1), 23-41. <https://doi.org/10.1177/153244000300300102>
- Turner, W. R., Nakamura, T., & Dinetti, M. (2004). Global urbanization and the separation of humans from nature. *Bioscience*, 54(6), 585-590. [https://doi.org/10.1641/0006-3568\(2004\)054\[0585:GUATSO\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2004)054[0585:GUATSO]2.0.CO;2)
- von Eschenbach, W. J. (2021). Transparency and the Black Box Problem: Why We Do Not Trust AI. *Philosophy & Technology*, 34(4), 1607–1622. <https://doi.org/10.1007/s13347-021-00477-0>
- Waldron, J. (2023). The Rule of Law. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/rule-of-law/>
- Wattles, J. (2023). Colonizing Mars could be dangerous and ridiculously expensive. Elon Musk wants to do it anyway. *CNN Business*. <https://edition.cnn.com/2020/09/08/tech/spacex-mars-profit-scn/index.html>
- Weijers, D. (2013). Intuitive Biases in Judgements about Thought Experiments: The Experience Machine Revisited. *Philosophical Writings*, 41(1):17-31.
- Weijers, D. (2014). Nozick's experience machine is dead, long live the experience machine! *Philosophical Psychology*, 27(4), 513-535. <https://doi.org/10.1080/09515089.2012.757889>
- Weijers, D., & Munn, N. (2024). Technology and Consumer Well-Being. *Fostering Consumer Well-Being: Theory, Evidence, and Policy*, 169-193.
- Westcott, K., Arbanas, J., Arkenberg, C., & Auxler, B. (2023, April 14). 2023 Digital media trends: Immersed and connected. *Deloitte Insights*. <https://www2.deloitte.com/us/en/insights/industry/technology/digital-media-trends-consumption-habits-survey/2023.html#read-the-digital-media-trends>
- William, B. (2017). Enlightenment. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/enlightenment/>
- Williams, D., Ducheneaut, N., Xiong, L., Zhang, Y., Yee, N., & Nickell, E. (2006). From tree house to barracks: The social life of guilds in World of Warcraft. *Games and culture*, 1(4), 338-361. <https://doi.org/10.1177/155541200629261>
- Wood, A. (2017, September 28). Getting Over It is a brutal new game from the maker of QWOP. *PC Gamer*. <https://www.pcgamer.com/getting-over-it-is-a-hellacious-new-game-from-the-maker-of-qwop/>
- Wu, C. (2024). Data privacy: From transparency to fairness. *Technology in Society*, 76, 102457. <https://doi.org/10.1016/j.techsoc.2024.102457>
- Zuboff, S. (2019). Surveillance Capitalism and the Challenge of Collective Action. *New Labor Forum*, 28(1), 10–29. <https://doi.org/10.1177/1095796018819461>