8-22-2024 1:00 PM

A Defense of Virtual Veridicalism

Yen-Tung Lee, *Western University*

Supervisor: Mendelovici, Angela, *The University of Western Ontario*
Co-Supervisor: Bourget, David, *The University of Western Ontario*
A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Philosophy

Abstract

Virtual reality is poised to be increasingly important in our lives. This dissertation investigates the philosophical foundations of virtual reality, probing the metaphysics and epistemology of perceptual experiences of virtual environments. Specifically, it asks 1) what there is in virtual reality and 2) how we perceive virtual things. It defends virtual veridicalism, the view that perceptual experiences in virtual reality are as veridical as ordinary experiences. The defense consists of six chapters. Chapter 1 explains why such questions need to be addressed. Chapter 2 develops a realist view of virtual objects and properties, arguing that virtual objects exist and instantiate properties that may or may not be instantiated by ordinary physical objects. Chapter 3 presents an argument in defense of virtual veridicalism. I argue that the common criterion whereby we consider ordinary experiences veridical justifies the veridicality of perceptual experiences of virtual objects and properties. Chapter 4 looks for a representationalist account of perceptual content that could lend theoretical support to my defense of virtual veridicalism. I argue that none of the currently prominent representationalist views succeed. Rather, Chapter 5 proposes a new account — role representationalism — to achieve the goal. This account claims that, despite some differences in perceptual contents between the virtual and ordinary physical domains, a fundamental aspect of contents remains constant in perceptual experience across both domains, and it is the constant aspect that grounds the veridicality of perceptual experiences of virtual objects and properties. Chapter 6 assesses Chalmers's (2017) defense of virtual veridicalism and argues that his defense cannot succeed unless presupposing role representationalism as the underlying theory of content of perceptual experiences. Finally, I conclude that 1) there are virtual objects and properties in virtual reality, and 2) we perceive these entities in a way that is as veridical as our perception of ordinary objects.

Keywords

Perceptual experience, virtual reality, virtual veridicalism, role representationalism, virtual digitalism.

Summary for Lay Audience

Virtual reality (VR) is becoming increasingly important in our lives. In VR, we perceive and interact with things that never happen in ordinary lives. For instance, in the game Beat Saber, players have to hit the virtual cubes that swiftly fly toward them with the virtual sabers in their hands. To reach a high score, they need to coordinate her movements with the rhythm of the flying cubes. Presumably, their hitting the cubes will succeed only when their visual experiences of the cubes are accurate, or more specifically, only when they see the cubes as approaching in the right position at the right time. In other words, for the players' movements to be successful, it seems to require that the cubes exist and fly in the game and that the players' experiences of the cubes' trajectory be accurate or, in philosophers' jargon, veridical. It might be hard to explain the success of their movements if they perceive the cubes in the wrong position or if the cubes are not really there.

The reflection above motivates two general philosophical questions of VR: 1) What is there in VR? 2) How do we perceive virtual things? The first question concerns the metaphysics of VR. The second question, on the other hand, focuses on the epistemic foundation of perception in VR. This dissertation aims to answer both questions. Its central thesis is virtual veridicalism, the view that perceptual experiences in VR are as veridical as ordinary experiences. First of all, I contend that virtual objects (e.g., the cubes) are real and instantiate perceptible properties (e.g., the cubes are flying). Second, I argue that our perceptual experiences of them are veridical, i.e., we perceive virtual objects as having the properties that they really have. Eventually, I conclude that 1) there are virtual objects and properties in VR, and 2) we perceive these entities in ways that are as veridical as our perception of ordinary objects.

To the memory of my grandfather

Acknowledgments

PhD is often a solitary voyage, yet the companionship of friends and colleagues made this journey a shared adventure.

Throughout this journey, I owe a deep debt of gratitude to my co-supervisors, Dr. David Bourget and Dr. Angela Mendelovici, whose steadfast enthusiasm and unconditional support have been my compass. Their commitment to excellence has been an endless wellspring of inspiration, and it has been a profound honor to work with them.

I extend my heartfelt thanks to the members of my prospectus and thesis examination committee: Dr. Chris Viger, Dr. Robert Stainton, Dr. Anthony Skelton, Dr. Terry Peters, and Dr. Adam Sennet. Our conversations were a source of enlightenment and their generous support left an indelible mark on me.

This dissertation flourished through the insights of scholars I encountered at various academic events. My gratitude goes to the audiences of my presentations at the "Consciousness and Reality" Graduate Conference 2022 (Chapter 6), the Annual Congress of the Canadian Philosophical Association (CPA) 2023 (Chapter 3), the 26th Annual Meeting of the Association for Scientific Study of Consciousness 2023 (ASSC26) (Chapters 4 & 5), the Consciousness Research Network (CoRN) 2023, (Chapter 6), and the "Perception in Games and Virtual Worlds" Workshop 2024 (Chapter 3). I also cherish the discussions with my fellow Taiwanese philosophers, Caleb, Chen-Wei, Shao-An, and Tony.

I am grateful for the financial support from the Western Graduate Research Scholarship (WGRS) during my first four years at Western University, and to the National Science and Technology Council of Taiwan for the 2023 National Science and Technology Council Taiwanese Overseas Pioneers Grants (TOP Grants) that funded the last year of my study. This dissertation would not have been possible without this support.

The daily interactions with my colleagues enriched my routine and propelled my research progress. I am thankful to my cohort, Donovan, Jaipreet, Varun, and Yichen, whose companionship throughout the pandemic and our monthly dissertation chapter reviews were invaluable for me to finesse my work.

Being part of the Philosophy Department and the Rotman Institute of Philosophy has been a privilege. I have made lasting friendships with (but not limited to) Adam, Amy, Angelica, Andrew,

Charlotte, Chris, Dmitry, Ed, Enzo, Georgia, Giovanni, Heather, Hershy, Jackie, Jacob, James, Jess, Jonathan, Kardelen, Mike, Niels, Owen, Pep, Rob, Robert, Sarah, Saul, Sid, Sijie, Todd, Wayne, Yusheng, and Zeruo. (If your name is missing, please forgive me — I promise to write a review for your work in the future!) The Rotman Institute provided a nurturing environment for my dissertation to blossom, and I will never forget the flavor of coffee there!

My Taiwanese/Chinese friends alleviated my homesickness. Special thanks to Yi-Hsuan for her generous culinary support and to Mingyan for sparking many intriguing ideas in my dissertation. I am also fortunate to have close friends like Shuwei, Xiaoyue, Xinyu, Yichen, Yifei, Yiying, Zhaoxuan, Zhuoyao, and Zili. Their companionship has been invaluable during my time in London.

I heartily thank my family for all their unwavering love and support: my parents, grandparents, siblings, aunts, uncles, cousins, my family-in-law, and all extended family members. Their encouragement allowed me to focus on my research without worry. Unfortunately, I no longer have a chance to show my dissertation to my beloved grandfather, who passed away in late 2023. This work is dedicated to him in memory of his love.

Lastly, my deepest gratitude goes to my wife, Tzu-Ying. Despite the physical distance imposed by the pandemic, her constant emotional support and encouragement have been my anchor. Tzu-Ying, thank you for everything we have experienced together, and I eagerly anticipate the future we will share.

Table of Contents

Chapter 1: Introduction

Virtual Reality (VR) is becoming increasingly important in our lives. Already, activities in many domains are accomplished through VR. In VR flight training simulations, trainees wearing VR devices learn to react appropriately to a myriad of situations they would encounter as if they were really piloting a plane. Through VR simulations, contemporary artworks generate aesthetic experiences previously unattainable in traditional genres. For example, the MIT art project *Tree* simulated a rainforest environment, allowing participants to embody a tree and "feel" their ecological relationship with other plants and animals.[1] VR also enables new e-sports. With an increasing number of VR games available globally on digital distribution services such as Steam, Origin, or Xbox Live Marketplace, the organization of tournaments, such as the VR Master League, has become a profitable industry on a global scale. These activities have achieved remarkable success thanks to the advancements of VR technology.

A deep understanding of how humans perceive and experience virtual environments is likely indispensable for the optimal exploitation of VR technology. For instance, a good understanding of the psychological processes of flight trainees can enhance the effectiveness of training programs by enabling the technology to suitably manipulate their senses during the training process. Likewise, a profound understanding of how aesthetic experiences are formed and influenced by virtual art can propel the progress of this medium. Similarly, game designers may be able to design more captivating and immersing games if armed with insights into why VR gaming experiences captivate players. These examples illustrate how a better understanding of our relationship to virtual realities can help unlock the potential of VR.

The evolution of VR not only enriches a diverse array of activities but also paints a compelling vision for the future. In October 2021, Facebook CEO Mark Zuckerberg unveiled his vision for a universal online VR social platform, the Metaverse. He predicted the possibility of building a cohesive society in this "new reality" by transcending the constraints of physical distance. The realization of such a society hinges on "a shared foundation of values and some understanding of the world and the problems that we all face together".[2] Without this common ground, chaos may ensue in the virtual world. For instance, if virtual objects such as non-fungible

[1] https://www.media.mit.edu/projects/tree/overview/
[2] https://www.brookings.edu/blog/techtank/2021/09/30/the-metachallenges-of-the-metaverse/

tokens are deemed real, we can formulate social rules based on their existence, regulating their possession and utility in ways analogous to laws governing tangible properties. If, however, virtual objects are considered fictional, regulations on them may lack a solid foundation. A shared epistemic foundation is necessary for a resilient virtual network.

Philosophy has great potential to contribute to a deeper exploration of the interrelationships between our daily lives and VR technology. For one, philosophy can provide a framework for the nature of various human experiences in virtual environments, such as perceptual experience, aesthetic experience, or even intersubjective experience. Knowing how humans experience virtual reality, VR developers would be able to improve the technology to create more vivid virtual environments. In addition, philosophy can identify ways to build a shared foundation that grounds the values and knowledge within the virtual environments. With such a foundation, the dream to create a distance-free global community would probably come true.

Moreover, the development of VR technology propels further philosophical reflections. To name a few, as VR technology creates perceptual environments that humans have never encountered before, philosophers could explore the nature of perception by examining how we perceive in virtual environments. If we perceive virtual things exactly in the same way as we perceive ordinary things, our perceptual content may not be determined by the particular object being perceived. Also, understanding people's reactions to alleged misconduct in the virtual realm would shed light on how philosophers approach reality and morality. If virtual rape (i.e., the rape of an avatar that represents a person in a virtual world) is considered morally wrong, we might be more inclined toward a view that the avatar is not just a "fictional copy" of ourselves but a genuinely agentive being whose rights should be protected as we protect our bodies in ordinary life.[3]

This dissertation addresses issues surrounding the reciprocal relationship between philosophy and VR technology. Specifically, this dissertation explores the epistemological and metaphysical foundations of VR. Drawing on the lessons from VR technology, I focus on human perception in virtual environments. Perception is the mind's first point of contact with the world. Although perception sometimes conveys inaccurate information, it is supposed to be systematically veridical in ordinary life. This dissertation extends this systematic veridicality to perception in VR, defending the view that many perceptual experiences in VR are veridical. This

[3] A concrete case of virtual rape will be extensively discussed in Section 1.1.2.

view is called *virtual veridicalism*. Of course, a full defense of such a view not only requires a detailed survey of perception but also demands a coherent picture that addresses philosophical issues regarding ontology, truth, meaning, knowledge, etc. As readers will see in the following chapters, these issues will play a significant role in my defense of virtual veridicalism.

In this dissertation, I adopt the approach of *technophilosophy*. This approach is proposed in Chalmers's book *Reality+: Virtual Worlds and the Problem of Philosophy*, published in 2022. According to Chalmers, the key to technophilosophy is the interaction between philosophy and technology. Specifically, this approach "1) asks philosophical questions about technology, and 2) uses technology to help answer traditional philosophical questions" (Chalmers, 2022: xviii). The reciprocal relationship between philosophical progress and technological development is central to this approach. This dissertation applies this approach exclusively to VR technology. I believe that the progress of philosophy and the development of VR technology benefit each other: VR technology finds its place in the modern world with the support of philosophical foundations while simultaneously providing new resources to reframe longstanding philosophical debates.

1.1. How Do Philosophy and Virtual Reality Benefit Each Other?

Before proceeding to the detailed objectives of this dissertation, this section aims to substantiate the reciprocal relationship between philosophy and VR technology and to point out how each benefits the other. I intend to show that these concrete examples serve as solid motivational grounds for my dissertation.

1.1.1. Why Does Philosophy Need Virtual Reality?

In his 2018 paper, Metzinger catalogs numerous intersecting points between VR technology and various areas of philosophy, hoping for more intensive interdisciplinary cooperation in the future. These intersections include philosophy of mind, epistemology, metaphysics, aesthetics, philosophy of technology, philosophy of action, social and political philosophy, etc. I extend Metzinger's catalog to address more concrete examples, showing that VR technology can help facilitate the progress of philosophy. Since this dissertation addresses the intersections specifically in philosophy of mind, epistemology, and metaphysics, this sub-section illustrates examples in these areas.

In philosophy of mind, perception is a sub-field highly relevant to VR technology. Perception is the mind's first contact with the external world. What appears to us in perceptual experiences is the basis of our knowledge of reality. But not all perceptual experiences offer justifiable grounds for knowledge. We might suffer from illusion or hallucination that is subjectively indistinguishable from genuine perception. If these kinds of errors are possible, how can we know anything about the world from experience? A traditional solution may fix perception at the external world, arguing that genuine perception is possible only if it is related to a physical object. However, such an approach may no longer target physical objects exclusively nowadays. Due to the advancement of VR technology, the once-sharp boundary between the physical and the virtual becomes blurred, leaving room for virtual objects also to be part of the external world. While playing *Pokémon GO*, for instance, we perceive a "pikachu". Via the aid of VR equipment, not only can the pikachu be perceived by many people from different perspectives, but it also persists "out there" through time. When a new pikachu appears somewhere, people gather together intending to catch it.[4] This case suggests that we can have genuine perceptual experiences of virtual objects. As this topic is a central concern of this dissertation, I will develop my view in the following chapters. My point here is that if perception of virtual objects and properties is possible, there seems to be room for philosophers to re-assess current theories of perception.

The development of VR also impacts epistemology. Philosophers have been struggling to escape from skepticism for centuries. According to an argument for skepticism, if we do not know we are not brains-in-a-vat, then we know nothing about the world. To secure our knowledge about the external world, many counterarguments against skepticism have been proposed, but even now, none of them is strong enough to dispel the threat of skepticism. The rise of VR technology brings such a traditional predicament back to life because it realizes metaphysically possible scenarios in the actual world. Those conceived in abstract thought experiments take place vividly in front of us nowadays, making relevant debates more complicated. In particular, we may be Neo in the movie *Matrix*, trying to clarify whether the world we inhabit from birth is really out there. We might not be so fortunate as Neo to realize the truth. We might live in a virtual environment forever and know nothing throughout our lives. Assuming that we know something about our world, I believe

[4] For video reports about people's gathering to catch new pokemons, please see < https://youtu.be/DNjpqm73GBY > or < https://youtu.be/QNRWZqMU16g >.

VR provides a great opportunity to re-assess how external knowledge is possible and, if possible, how to secure it.

VR also helps reshape metaphysical inquiry. What is real? Throughout history, this question has been one of the most commonly asked metaphysical questions. But such an old question never comes to an end. Technology creates novel domains of entities that are also subject to this question. For instance, we can ask whether the airplane seen in VR flight training is real or not; we can ask whether the tree we embody during a demonstration of the MIT art project is real or not; we can also ask whether the weapons we use in a virtual 5v5 first-person shooting game are real or not. In general, we can ask whether entities in the virtual domain are real or not. The answer to this question affects how we probe the issues discussed above. Presumably, if virtual entities are real, we could naturally build knowledge about virtual reality through our perception of these "real" entities. Without their existence, we seem unable to perceive anything in VR, and our experiences in virtual environments might become completely illusory. The following chapters will show the plausibility of this idea. To make the inquiry possible for now, we should better not deny the existence of virtual entities at the very beginning.

Although there are many more questions to be asked, thanks to the advancements of VR technology, I believe the discussions above are sufficient for us to appreciate the contribution of VR technology to the progress of philosophy. As many core issues in philosophy would be reshaped by the rise of VR, I believe philosophers could re-examine their focus by lessons learned from VR. As we will soon see, this dissertation is primarily motivated by such an idea.

1.1.2. Why Does Virtual Reality Need Philosophy?

Having shown how VR technology can contribute to philosophy, this sub-section demonstrates cases the other way around, i.e., cases in which philosophy helps enhance the future application of VR. I aim to show some potential contributions of philosophy to VR.

A crucial import of VR technology to our lives lies in its capacity to create extremely vivid, immersive virtual environments. Users usually feel as if they were present in the virtual environment. They usually take the objects appearing in their surroundings as real. They also experience events that take place in VR as really happening. In this respect, we are in a great position to ask why this is the case: Why do we experience a sense of presence in the virtual environment? Why do we take virtual objects as real? And why do we experience a sense of reality

in VR? Although a common response tends to dismiss them for the presupposition that VR is a grand illusion by its nature (e.g., Slater, 2009), I think that an adequate answer to these questions would be much more interesting and profound. Presumably, we experience the sense of presence and the sense of reality in an ordinary environment because the environment and its objects and events are real. By analogy, perhaps we experience the sense of presence in a virtual environment because the environment is real; perhaps we take virtual objects as real because they are real; and perhaps we experience events in the virtual environment as really happening because these events really happen in the virtual environment. In fact, this is a consequence of the main thesis of this dissertation: Many perceptual experiences of virtual things are veridical partly because these virtual things are real.

One might wonder what implications this consequence has for VR. I want to suggest that it helps us envision our future and the role of VR in it. If virtual realities are real, our dream (or at least Zuckerberg's) to build a cohesive society in the Metaverse becomes more realistic. Since the maintenance of a cohesive society is constrained by how its members are ruled and regulated, and for rules and regulations to be effective for everyone, they must be anchored on real objects whose existence is widely accepted. For instance, the existence of personal digital properties is the basis of the relevant regulations, such as the European Union's Markets in Crypto Assets Regulation (MiCAR). Otherwise, nothing would be subject to these regulations.

In addition, our experience also serves as an indispensable component of the maintenance of the Metaverse society. In late 2023, UK police initiated an investigation of an alleged crime of "virtual rape". Using the Metaverse, a 16-year-old girl's virtual avatar is said to have been sexually attacked by a group of online strangers. Although her physical body remained unharmed as there was no physical assault in person, it is reported that "this child experienced psychological trauma similar to that of someone who has been physically raped", and that "there is an emotional and psychological impact on the victim that is longer term than any physical injuries".[5] In fact, prosecuting such a case under existing UK laws is practically impossible, since sexual assault is literally defined as non-consensual "physical contact" in a sexual manner. However, we would probably be inclined to consider the case as a sexual assault and to impose a sanction upon the aggressor. After all, we take the victim's experience as sufficiently genuine and veridical when

[5] Online resource is adapted from < https://www.ndtv.com/world-news/uk-girl-virtually-gang-raped-metaverse-virtual-reality-game-cops-begin-probe-meta-4790442 >.

and after her being attacked. If the experience were a mere illusion, the event would probably not affect the victim's mentality for so long as if she were physically attacked, and the relevant treatments would become unmotivated. To maintain the cohesion of the Metaverse, we would better find a way to resolve such a problem.

More issues arise while philosophizing the victim's experience in this case. Firstly, physical contact seems unnecessary for sexual behaviors. As virtual sex has become much easier to realize in virtual contexts nowadays, it might no longer be appropriate to restrict sexual behaviors to those involving physical contact. In this sense, sexual attacks might also happen without physical contact. Secondly, one's avatar might also count as (part of) one's body. Our self seems to inhabit our skull and skin. But recent studies suggest that our self can extend beyond biological boundaries and incorporate other objects into "one's body" (see de Vignemont, 2018). So, it may be possible for one to incorporate an avatar as part of one's body in an immersive virtual context. If so, a part of the victim's body was really assaulted in the event. Thirdly, and most fundamentally, the avatar seems a real object subjected to causal interaction with others. It seems that virtual rape is possible only if there *is* something to be attacked in the Metaverse. If the avatar and the sexual assault were merely fictional rather than real, the negative impact on the victim would not be so overwhelming. Therefore, philosophical reflections are needed for the applied cases of VR to resolve practical issues in the Metaverse and keep our virtual society safe.

The contribution of philosophical reflection to VR application may become widespread in enhancing social engagement or emotional attachment. VR has been employed to promote people's well-being for years. It is reported that elderlies suffering from social isolation feel less lonely while interacting with others' avatars with the aid of VR social platforms.[6] During the COVID-19 pandemic, people under the social distancing policy benefited from engaging in social activities in virtual environments without leaving their houses, and such activities significantly reduced their social anxiety (Kenyon et al., 2023). These cases indicate that VR has great potential to satisfy people's psychological need for social engagement even without in-person contact, triggering concern about whether the human need for emotional attachment could be satisfied by interactions with non-person avatars. If the answer is positive, a huge amount of philosophical and psychological literature in this area might be undermined, since the involvement of other persons

[6] For some news reports, see also < https://www.forbes.com/sites/solrogers/2020/02/26/how-virtual-reality-is-benefiting-seniors/?sh=4145bf331485 >.

is deemed necessary for satisfying the need. But if the answer is negative, it seems that the avatars could take over the emotional role played by real persons, or even that avatars could become real persons.

A recent case appropriately embodies the potentiality of avatars to be real persons. This year, a Spanish artist, Alicia Framis, is going to marry an interactive, holographic virtual avatar, AiLex, embedded with an artificial intelligence trained using profiles from Framis's previous relationships.[7] In their relationship, AiLex can do chores that look like what a normal person usually does. He (or it) can chat with Framis. He can have a virtual meal with Framis. He can also wash virtual dishes. Although AiLex is not touchable in the physical sense, it is reported that the couple engages in an intense romantic relationship, and AiLex "fulfills all Framis's emotional needs". In an interview, Framis expressed her vision of a posthuman society in which "humans will marry and maintain relationships with holograms, avatars, robots, and so on". She believes that virtual partners could accompany one's life by providing long-term mental support, including empathy, intimacy, and emotional attachment. Eventually, says Framis, "love with holograms is an inevitable reality".

AiLex's case suggests that virtual partners can do what can be done only by real persons. Interpreting it at face value, the case aligns with the thought that virtual avatars may become real persons. Of course, they are not physically touchable. But they seem to be conscious and have a fruitful mentality according to which we treat them in the same way we treat real persons. In this sense, virtual partners can be real. Unlike the case of virtual rape, AiLex is not an extension of a real body. Instead, he is real by himself, though his body is not commonsensically physical. Also, AiLex's case is one in which VR contributes to humans' well-being. Having a relationship with a virtual partner is not a fictional play but a part of real life. Framis need not "pretend" to love AiLex. Rather, she "loves" him genuinely. Since such love is real, Framis is able to be emotionally attached to and share life with AiLex. Reflection on AiLex's identity also sheds light on the future of VR. Building on AiLex's case, VR developers can incorporate AiLex and a social robot in the future, so that Framis can not only live with AiLex but also physically interact with him. Given

[7] For pictures, videos, and more information of the hybrid couple, please see < https://www.euronews.com/culture/2024/01/03/meet-spanish-artist-alicia-framis-the-first-woman-to-marry-a-hologram > and < https://www.aliciaframis.com/work/154/hybrid-couple >.

the human need for physical contact with others, AiLex's case could serve as a starting point to facilitate our future well-being.

Admittedly, the descriptions of the above cases gloss over many subtleties. However, I hope this sub-section will motivate those interested in VR to recognize the contribution of philosophy to this technology. As a first attempt, this dissertation addresses some of the most fundamental topics in philosophy of virtual reality. The following chapters will focus on the metaphysics and epistemology of perceptual experience in VR. I believe my contributions to these topics will serve as the theoretical basis to resolve other issues around the application of VR.

1.2. Aims, Terminology, and Objectives

This dissertation's main thesis is *virtual veridicalism*, the view that many perceptual experiences of virtual things are veridical. To defend this thesis, I address two foundational questions: *What is there in virtual reality? How do we perceive virtual things?*

First, some terminology. An *experience* is a mental state that has phenomenal characters or what-it-is-likeness (Nagel, 1974). *Perceptual experiences* are experiences in sense modalities. For convenience, by experience, I mean perceptual experience unless the context indicates otherwise. *Veridical experiences* are, to a first approximation, experiences that represent the environment *as it really is*, and *falsidical experiences* otherwise.

By *VR experiences*, I mean experiences generated through the use of VR technology. I take the current best VR technology as the paradigm case of VR technology. This technology consists of a VR suit (e.g., the Teslasuit) and a headset that synchronously manipulates the wearer's exteroception.

It could turn out that the environment we inhabit from birth is itself a simulation (Bostrom, 2003; Chalmers, 2022). If so, all ordinary experiences are literally VR experiences. Call the virtual environments that we inhabit from birth *natural VR*. This dissertation does not address natural VR but only *invented VR*, i.e., the virtual environments we create through VR technology. I assume that we do not inhabit a natural VR. I also assume that ordinary, non-VR experiences are typically veridical.

By *virtual world*, I mean a shared, coherent simulated scenario. A *scenario* is a 3D environment expanded from one's first-person perspective, and a *simulated* scenario is a scenario simulated using computer programs. A scenario is *coherent* when its simulated space is fully

continuous. A scenario is *shared* when different subjects can get access to things in the scenario from different perspectives at different times, so the presence of these things is independent of any individual subject's experience. Crucially, a virtual world is *not* a counterfactual world. A virtual world is a scenario in the actual world. Simulation technology will create virtual worlds that may or may not *completely* screen off non-virtual things and then appear the same as non-VR environments. I call a virtual world that does so a *perfect virtual world* and a virtual world that does not a *mixed world*. Unless specified, I take a perfect virtual world as the paradigm case when discussing a virtual world.

This dissertation has three parts. The first part (Chapter 2) addresses the metaphysics of virtual things, answering the question of *what there is in VR*? I believe this question has to be answered before we develop an account of the veridicality of VR experiences. Presumably, since perceptual experiences can be veridical only if there are objects and properties being perceived, the instantiation of these properties and the existence of these objects are pre-conditions for the relevant experiences to be veridical. This part puts forward a realist view of virtual objects and properties. This realist view serves as the metaphysical basis for my arguments for virtual veridicalism in subsequent chapters.

The second part (Chapters 3–5) aims to explain *how we perceive virtual things* by defending virtual veridicalism. Having established the reality of virtual objects and properties in Chapter 2, I proceed to argue that our perceptual experiences of these objects and properties are plausibly veridical. This part presents two independent but mutually supporting arguments for virtual veridicalism. The first argument (in Chapter 3) defends virtual veridicalism by analyzing what it takes for a perceptual experience to be veridical. The second argument (in Chapters 4 and 5) defends virtual veridicalism by searching for an appropriate account of the contents of perceptual experience that not only preserves our intuition of the veridicality of ordinary experiences but also grounds the veridicality of VR experiences. Chapter 4 looks for such an account among currently prominent representationalist accounts, arguing that none of them fits. Chapter 5 proposes a new account that coheres well with the picture of virtual veridicalism. I conclude that virtual veridicalism is a plausible view of our perceptual experiences of virtual things.

The third part (Chapter 6) compares Chalmers's (2017) defense of virtual veridicalism with my defense. I contend that an internal inconsistency in Chalmers's account results in unacceptable

theoretical consequences. Identifying these tensions, I argue that the best resolution for his defense is to presuppose the account proposed in Chapter 5.

1.3. Chapter Overview

This section summarizes each chapter of the dissertation.

Chapter 2 probes the issues of virtual ontology, arguing that virtual objects (i.e., objects that we commonly talk about in a virtual world) exist and instantiate properties that may or may not be instantiated by ordinary, non-virtual objects. I argue against virtual irrealism and defend a version of *weak virtual digitalism* according to which virtual objects are grounded in data structures that are physically realized. For instance, an abstract array of bits 011001 is realized by a computer's underlying array of concrete memory cells with the corresponding array of voltages (i.e., low, high, high, low, low, high), and the abstract arrays further constitute high-level 3D graphics such as a "virtual dolphin". Also, I analyze the nature of virtual properties (i.e., properties that we commonly ascribe to virtual objects in a virtual world). I argue that some virtual properties are never instantiated by ordinary objects (e.g., *being a virtual kitten*), while others are identical to those instantiated by ordinary objects (e.g., *being a library*). Therefore, I conclude that virtual objects are real and instantiate virtual properties.

Chapter 3 argues for *virtual veridicalism,* the view that many VR experiences are veridical. I defend this view with a three-stage argument centered on perceptual variation, the phenomenon in which different perceptual experiences are qualitatively identical while being generated by different properties under different circumstances. The first stage argues that perceptual variation can occur in color perception without involving misperception. In particular, Color-inverted Earth (Block, 1990, 1996) is possible, and neither Nonverts nor Inverts suffer from systematic misperception. The second stage extends the argument to space perception. I argue that since experiential space (e.g., the space we perceive the external reality as having) is not identical to physical space, it is possible for individuals to perceive distinct physical spaces as having the same experiential space without suffering from systematic misperception. The final stage of the argument builds on the previous two stages to argue that perceptual variation without misperception in color and space perception can occur across virtual and ordinary environments. Assuming that ordinary experiences are largely veridical, many perceptual experiences in virtual reality are also veridical.

Chapter 4 argues that none of the current major accounts of phenomenal content coheres with virtual veridicalism. In particular, I reject physical Russellian representationalism, dispositionalist Russellian representationalism, and Fregean representationalism. First, since physical Russellian representationalism identifies what an experience represents with physical properties, it is impossible for an experience to represent virtual properties. Thus, the account predicts that virtual veridicalism is false: No VR experience is veridical. Second, I argue that dispositionalist Russellian representationalism faces an inherent problem, making the account so unstable that it cannot serve as an underlying account of content that coheres with virtual veridicalism. Third, I argue that Fregean representationalism is good enough to cohere with virtual veridicalism in perfect virtual worlds. However, such coherence is too restricted to be generalized to mixed worlds. As virtual veridicalism should accommodate the veridicality of perceptual experiences in both perfect virtual worlds and mixed worlds, Fregean representationalism ultimately fails to cohere with this view.

Chapter 5 proposes a new account of phenomenal content in the representationalist framework. I call the account *role representationalism*. Representationalism about phenomenal experience is roughly the view that every phenomenal character of an experience is determined by its representational content. Role representationalism adds that the phenomenal content (i.e., representational content of an experience that supervenes on its phenomenal character) of an experience attributes the higher-order property of *having a realizer that normally causes the experience type in the subject* under certain perceptual circumstances. I argue that this account perfectly coheres with virtual veridicalism in that the role representationalist framework accounts for perceptual variation without misperception across virtual and ordinary environments. In effect, reddish experience can be veridical in both virtual and ordinary environments because it represents the higher-order property of *having a realizer that normally causes reddish experience in the subject*, which can be realized by a surface reflectance profile in an ordinary environment or by a virtual property in a virtual world.

Chapter 6 assesses Chalmers's (2017) defense of virtual veridicalism. First, I argue that the key notion of Chalmers's defense, i.e., cognitive orientation, is ambiguous, and the interpretation that captures what Chalmers means in his defense is insufficient to defend virtual veridicalism. Specifically, his defense requires an underlying account of phenomenal content. Based on the conclusion of Chapters 4 and 5, I argue that Chalmers's defense of virtual veridicalism cannot

succeed unless it presupposes role representationalism as the underlying account. Therefore, Chalmers's defense of virtual veridicalism is incomplete without my own defense of the view.

Chapter 2: On Virtual Objects and Virtual Properties

Through VR devices, we interact with many things simulated by underlying computational processes. We may see a "virtual apple"; we may use a "virtual weapon" to win a "virtual battle"; we may also stretch our "virtual arm" to touch a "virtual fairy". Apples, weapons, battles, and fairies are things in virtual worlds that we can talk about. Let's call these things *virtual objects*.[1] In addition to virtual objects, we can talk about their characteristics: We may describe the virtual apple as being "red"; we may say the virtual weapon is "useless" in defending the "hard" battle; we may also realize that our virtual arm is not "long" enough to reach the "flying" fairy. These characteristics seem to be the properties that we take virtual objects to have. Call them *virtual properties*.[2] Through discourse about virtual objects and virtual properties, VR users communicate and understand each other. For instance, in defending the team's stronghold in an e-sport game, teammates swiftly share information regarding the enemies' equipment and location by ascribing virtual objects and properties to the gaming environment. The defense cannot succeed without effective communication via teammates' understanding of what others are talking about.

Although virtual objects and properties are easily talked about through ordinary language, their metaphysics remains unclear. The fact that we ascribe virtual objects and properties to a virtual world in ordinary language does not imply that virtual objects are real or that virtual properties are instantiated. Virtual objects may be real objects or fictional objects, and virtual properties may or may not be really instantiated. So, the question is: Are virtual objects real, and are virtual properties instantiated?

As a preliminary to my later defense of virtual veridicalism, this chapter aims to clarify the metaphysics of virtual worlds. I argue for two claims: (1) Virtual objects are real, and (2) they instantiate virtual properties. For (1), Section 2.1 and Section 2.2 defend a realist view, building on Chalmers' (2017, 2019) virtual digitalism and addressing criticisms directed against this view. Section 2.3 argues for (2). Distinguishing between two kinds of properties, *categorical properties* and *role properties*, it argues that many virtual properties are either categorical properties or role properties, and they are instantiated by virtual objects.

[1] For convenience, I take virtual objects as a sort of concrete objects.
[2] I use "property" as a generic term that covers anything that is predicable, exemplifiable, and viewable as universals, unless the context suggests otherwise (see Orilia & Paoletti, 2020).

2.1. Virtual Digitalism

Are virtual objects real? More specifically, do things in virtual worlds that we "talk about" exist? I consider two possible answers to this question:

Virtual object realism: Some virtual objects really exist.
Virtual object irrealism: No virtual objects really exist.

Realism admits the metaphysical status of virtual objects, whereas irrealism does not. There are different versions of them. In his 2017 paper titled "The Virtual and the Real", Chalmers defends *virtual digitalism*, a version of realism that consists of two claims:

VD1 Virtual objects really exist and are digital objects.
VD2 Events in virtual worlds are largely digital events that really take place. (Chalmers, 2017: 311) [3]

In this chapter, I focus on the arguments for and against VD1, which concerns the existence and the grounds (i.e., the metaphysical explanations) [4] of virtual objects. VD2 is irrelevant for my purposes because it does not concern objects. For convenience, by virtual digitalism, I mean *virtual digitalism about virtual objects*, i.e., VD1. Now, let's overview Chalmers's arguments for VD1.

VD1 consists of two conjuncts. The first conjunct (i.e., virtual objects really exist) is exactly virtual object realism, so the latter is true if VD1 is true. The second conjunct is a metaphysical claim about the identity of virtual objects and digital objects. Digital objects, to a first approximation, are data structures "grounded in computational processes which are themselves grounded in physical processes on one or more computers" (*Ibid.*: 317). To a second

[3] Chalmers's definition of "virtual worlds" largely overlaps with mine, according to which a virtual world is a shared, coherent simulated scenario (see Chapter 1). By a virtual world he means "an interactive computer-generated environment, of the sort that we (seem to) inhabit when using virtual reality" (*Ibid.*: 314). The nuance between our definitions does not matter to this chapter, so I assume we mean the same when using the term.
[4] In the literature, some philosophers claim that grounding just is metaphysical explanation (e.g., Dasgupta, 2014), while some deny this claim, arguing that grounding provides support to a metaphysical explanation (e.g., Audi, 2012). Even some deny a necessary connection between grounding and metaphysical explanations (e.g., Maurin, 2019). As the nuance does not affect my purposes, I assume that metaphysical grounding just is metaphysical explanation for convenience of discussion.

approximation, however, digital objects are higher-level entities "constituted by these data structures" (*Ibid.*: 317). The first approximation identifies digital objects with data structures, while the second approximation commits to a weaker, grounding relation between them. Either way, digital objects are grounded in data structures.

Chalmers's usage of "data structures" is idiosyncratic. By a data structure, he means a *physical data structure*, which is "a *realization* of the abstract data structures specified in computer programs and the like" (Chalmers, 2019: 460; original emphasis). A realization of abstract data has to occur in concrete computational systems. For instance, an array of bits 011001 is realized by a particular computer's underlying physical data structure, such as an array of concrete memory cells with the corresponding array of voltages (low, high, high, low, low, high). In this case, the array of voltages is the physical data structure that realizes the array of bits, which is the abstract data structure.[5] Moreover, high-level data structures such as serial numbers or even graphics are "implemented (after compilation and/or interpretation) in certain low-level structures involving bits, which are themselves physically realized when a program is executed" (*Ibid.*: 462–3). One might also impose more mathematical structure upon the abstract array so that it is structurally isomorphic to the original physical data structure. No matter how complex the structure is, it is plain that "the physical data structure that realizes the high-level data structure is the array of concrete bits that realize the formal bits that implement the data structure" (*Ibid.*: 463).

To continue, Chalmers (2017) offers two arguments to defend virtual digitalism. One regards the causal power of virtual objects, and the other concerns their role in perception in virtual worlds. First, for the argument from causal power, Chalmers argues that the causal powers that virtual objects have are constituted by the causal powers of digital objects, so it is plausible that virtual objects are digital objects. Presumably, virtual objects causally interact. In a video game, I control my virtual avatar to kill a virtual zombie, suggesting a causal relationship between the avatar and the zombie. My avatar is constituted by a certain data structure, and so is the zombie. Also, their causal interaction is always accompanied by an interaction between their underlying data structures. Whenever the avatar and the zombie interact, their data structures interact. According to Chalmers, such co-existence of the causal interaction between virtual objects and

[5] Admittedly, many critics argue that Chalmers's non-standard usage of "data structure" results in a confusion in the relevant debate (e.g., Ludlow, 2019; Beisbart, 2019). But for the present purposes, I simply follow Chalmers's terminology and leave open to the correct usage of this term.

that between data structures suggests that the former just is the latter. Thus, the argument from causal power is formulated as follows:

AC1 Virtual objects have certain causal powers (to affect other virtual objects, to affect users, and so on).

AC2 Digital objects really have those causal powers (and nothing else does).

AC3 Therefore, virtual objects are digital objects. (*Ibid.*: 318)

A second argument, the argument from perception, appeals to what VR users perceive in a virtual world. Since the digital objects generated by computer programs are causally responsible for the users' perceptual experiences, as Chalmers contends, the virtual objects they perceive are digital objects.

AP1 When using virtual reality, we perceive (only) virtual objects.

AP2 The objects we perceive are the causal basis of our perceptual experiences.

AP3 When using virtual reality, the causal bases of our perceptual experiences are digital objects.

AP4 Therefore, virtual objects are digital objects. (*Ibid.*: 318)

AP1 is intuitively plausible. AP2 is widely accepted by philosophers of perception. AP3 seems empirically correct, as the data structure that constitutes a virtual object that one perceives is causally responsible for generating one's visual experience.

These two arguments are not knockdown arguments against rival views. But I am not going to defend each premise. For present purposes, we can safely suppose that these arguments provide *prima facie* justification for VD1.

However, VD1 seems too strong to be true. Specifically, the identity between virtual objects with digital objects is challenged by the case of cross-play games (a.k.a. cross-platform play), a sort of video game in which multiple players can play together using distinct gaming platforms:

> "Suppose that *Frisbee* is a cross-play-supporting VR application of a virtual environment wherein multiple users can play virtually throw [*sic*.] around a virtual frisbee. Suppose further that Neil, on his HTC Vive in 32-bit mode, and Nathan, using his Oculus Rift (which only supports 64-bit mode), are playing *Frisbee* together over a peer-to-peer network (i.e. there is no mediating server that controls the game). In such a case, there is a single virtual object that Neil and Nathan are both engaged with — the virtual frisbee. However, as the local hardware and system architecture of their respective VR set-ups are radically different and, since there is no common server, there is no single digital object that both Neil and Nathan are engaged with." (McDonnell & Wildman 2019: 379).

The cross-play problem parallels the argument from multiple realizability against the identity between high-level and low-level objects. The gist is that distinct digital objects can realize a single virtual object. Specifically, a set of 32-bit or 64-bit data structures can realize the same virtual frisbee. The 32-bit data structure is distinct from the 64-bit data structure, so the virtual object cannot be identified with any of them. Therefore, VD1 is false.

A simple resolution of the cross-play problem is, as Chalmers does in the response paper titled "The Virtual as the Digital" (2019), to insist that the cross-play case involves only one digital object identical to the virtual frisbee. He takes the second approximation of digital objects, according to which digital objects are higher-level entities constituted by data structures, so "[the cross-play case] is a case where a single virtual object [= a single digital object] is realized jointly by two data structures" (*Ibid*.: 471). In this way, the cross-play problem does not arise because McDonnell & Wildman mistakenly identify digital objects with data structures.

While being coherent, I do not think this response succeeds because it makes the nature of digital objects mysterious. Firstly, the individuation condition of digital objects becomes unclear. A digital object cannot be individuated by data structures, as it is multiply realizable by the latter. Nor can it be individuated by virtual objects, since the nature of digital objects is supposed to explain the nature of virtual objects. I think an ideal solution would be to identify digital objects with a certain level of abstraction found in high-level programming languages. But nuances still need to be settled to demystify what a digital object is. Secondly, it seems *ad hoc* to identify virtual objects with mysterious digital objects without an account of the latter. As the nature of digital objects is unknown, it is useless to explain the nature of virtual objects as VD1 does.

To avoid the mysteries introduced with digital objects, VD1 can be reformulated as follows:

VD1* Virtual objects really exist and are metaphysically grounded on *data structures*.

And according to the two approximations mentioned above, VD1* breaks down into two versions:

VD1*(S) Virtual objects really exist and *are identical to* (really existing) data structures.

VD1*(W) Virtual objects really exist and *are metaphysically grounded on* (really existing) data structures.

Call virtual digitalism with VD1*(S) *strong virtual digitalism*, and call virtual digitalism with VD1*(W) *weak virtual digitalism*.[6] The cross-play problem offers a strong reason to reject VD1*(S). If virtual objects are digital objects, and digital objects are data structures, then two distinct data structures cannot realize the same virtual object. Yet, virtual digitalists need not be strong virtual digitalists. They can retreat to weak virtual digitalism and commit to VD1*(W), allowing virtual objects to be distinct entities from data structures. For weak virtual digitalists, the cross-play problem is simply unmotivated because the metaphysical grounding of virtual objects on data structures is compatible with the multiple realizability of virtual objects.

Due to the retreat to weak virtual digitalism, Chalmers's initial arguments (i.e., the argument from causal powers and the argument from perception) should be revised accordingly. First, the argument from causal power (weak version):

AC1*(W) Virtual objects have certain causal powers (to affect other virtual objects, to affect users, and so on).

AC2*(W) All causal powers the virtual objects have are dependent upon the causal powers *data structures* have (and nothing else does).

AC3*(W) Therefore, virtual objects are dependent upon *data structures*.

[6] Likewise, VD2 can also be distinguished into two versions that correspond to VD1*(S) and VD1*(W):
VD2(S): Events in virtual worlds *are identical to* digital events.
VD2(W): Events in virtual worlds *are metaphysically grounded on* digital events.

Also, the argument from perception (weak version):

AP1*(W) When using virtual reality, we perceive (only) virtual objects.
AP2*(W) The objects we perceive are the causal basis of our perceptual experiences.
AP3*(W) When using virtual reality, all the causal bases of our perceptual experiences are dependent upon the causal powers of *data structures*.
AP4*(W) Therefore, virtual objects are dependent upon *data structures*.

Apart from the premises in which digital objects are replaced with data structure, the premises being weakened in these arguments are AC2*(W) and AP3*(W), respectively. The truth of these premises coincides with an assumption that the dependency of causal powers suffices for the dependency of the relevant objects: If any virtual objects exist and have causal powers, their causal powers must depend upon the causal powers of data structures that ground those virtual objects. I think the assumption is plausible:

> "Corresponding to each avatar in [the video game] *Second Life*, there is a data structure on the *Second Life* servers. [...] When I see an avatar, it is this data structure that brings about my perception. What I perceive directly reflects the properties of this data structure: the perceived location of the avatar reflects one property of the data structure, while the perceived size, color, and so on reflect other properties. When my avatar interacts with a coin, the two data structures are interacting. Whenever two virtual objects interact in *Second Life*, there is a corresponding interaction among data structures. Data structures are causally active on real computers in the real world; the virtual world of *Second Life* is largely constituted by causal interaction among these data structures." (Chalmers, 2017: 317)

According to this passage, virtual objects interact by virtue of the interactions between underlying data structures, and the causal powers of virtual objects are grounded in the causal powers of data structures. Without the interaction of the grounds, there cannot be interaction of the grounded. Therefore, AC2*(W) and AP3*(W) are true. The weak version of both arguments justifies weak virtual digitalism.

However, another problem challenges weak virtual digitalism. McDonnell & Wildman (2019) adapt Kim's (1998, 2005) causal exclusion problem against non-reductive physicalism in special sciences. The problem consists in the inconsistency between four intuitive claims:

- *Non-reductive Physicalism*: Higher-level entities (i.e., entities studied in special sciences) non-reductively supervene on physical entities.
- *High-level Causation*: Higher-level entities cause physical entities.
- *Physical Causal Closure*: Every physical effect has a causally sufficient physical cause.
- *Exclusion Principle*: Causal effects are not causally overdetermined.

Take folk psychology as an example. *Non-reductive Physicalism* about the mental is the view that the mental non-reductively supervenes on the physical, e.g., the desire to raise my hand (*M*) supervenes on a certain neural state (*N*). *High-level Causation* states that *M* causes my behavior to raise my hand (*B*). *Physical Causal Closure* is widely accepted because we know from physical sciences that any physical effect must have a physical cause. So, it seems that *B* is caused by two distinct states *M* and *N*. But since the *Exclusion Principle* is presumably true, *B* can be caused by only one state, i.e., either *M* or *N*, but not both. The consequence is either that the causal efficacy of *M* is reduced to the causal efficacy of *N*, or that *M* does not really cause *B*. The first horn rejects *Non-reductive Physicalism* and retreats to a reductive view, while the second horn gives up *High-level Causation*, taking *M* to be causally inefficacious. Since neither horn is acceptable to non-reductive physicalists, the exclusion problem seriously challenges this position.

Similarly, McDonnell & Wildman (2019) argue that if we assume that (1) virtual objects have causal efficacy and that (2) the domain of data structure is causally closed, then either (3) virtual objects are reduced to data structures or (4) they are causally inefficacious. (1) is the central tenet of the weak version of the arguments from causal power and from perception, so weak virtual digitalists must accept (1). Since the causal interaction between data structures can be fully explained within the domain of data structure, (2) is also plausible. (1) and (2) together imply that "the digital process does all the causal work, leaving none for the virtual processes to do" (*Ibid.*: 384). In this sense, weak virtual digitalists are forced to either reduce virtual objects' causal powers to those of data structures or admit that virtual objects are causally inefficacious. The first horn is unacceptable to weak virtual digitalists, as it returns to strong virtual digitalism by identifying

virtual objects with data structures. But the second consequence would invalidate the arguments for weak virtual digitalism. If AC1*(W) and the conjunction of AP1*(W) and AP2*(W) are unmotivated, the initial justification of weak virtual digitalism is also called into question.

A resolution to the causal exclusion problem is to reject the *Exclusion Principle*. In the response paper, Chalmers (2019) argues that the exclusion problem fails "especially when there is a grounding connection between the lower-level and higher-level objects and properties in question…

> … If biological properties are grounded in [fundamentally] physical properties, then causal work by the [fundamentally] physical does not exclude causal work by the biological. Similarly, if virtual objects and properties are grounded in [data structures], then causal work by [data structures] does not exclude causal work by the virtual." (*Ibid.*: 472)

I believe Chalmers's response is on the right track, but he does not further argue for this passage in the response paper. I am going to elucidate his reasoning. His response is based on a general solution to the causal exclusion problem. The problem originally targets non-reductive physicalism in special sciences, including folk psychology, biological science, social science, and digital science (if any). Weak virtual digitalism can be taken as an instance, as it is also a non-reductive physicalist view of certain higher-level entities, namely, virtual objects. If the exclusion problem is blocked at the general level, such a resolution applies to each relevant view in all special sciences.

Recall the case in folk psychology. My behavior of raising my hand, *B*, may be caused by my desire to raise my hand (according to *High-level Causation*), *M*, and my neural state on which *M* supervenes (according to *Physical Causal Closure*), *N*. However, the *Exclusion Principle* requires that *B* be caused either by *M* or *N* but not both; otherwise, *B* will be causally overdetermined. To reject the *Exclusion Principle*, a strategy is to argue that *B* is not causally overdetermined even if caused by *M* and *N* at once because of a sufficiently strong metaphysical relation between *M* and *N*.[7] For instance, Yablo (1992) argues that overdetermination is resolved

[7] It is a terminological issue whether to characterize the metaphysical relation between *M* and *N* as resulting in a "benign" overdetermination or an elimination of overdetermination. I arbitrarily choose the second way, i.e., I take that the sufficiently strong metaphysical relation between *M* and *N* makes *B* no longer causally overdetermined. Afterall, this strategy argues that the case in which *B* is caused by *M* and *N* is *importantly different* from those we commonly take as problematic cases, say, a case in which the death of a firing squad victim is causally overdetermined by the bullets shot by members from the firing squad.

when *M* and *N* bear a determinable-determinate relation. The fact that both *M* and *N* are the causes of *B* does not result in overdetermination because if *N metaphysically determines M*, then it is metaphysically necessary that anything that is *N* is also *M*. Given that *N* causes *B*, *M* determined by *N* can also be cited in the causal explanation of *B* without making *B* overdetermined. Hence, *M* and *N* are not causally competitive. It is just a matter of how specific the answer we want to talk about the causal relation in question.

Analogously, the causal exclusion problem with respect to virtual objects can also be resolved in the way that the grounding relation between virtual objects and data structure defuses the threat of overdetermination. By definition, grounding is an irreflexive, asymmetric, and transitive relation of non-causal determination. The grounding relation is supposed to be under the subject of Yablo's argument. In this way, the threat of overdetermination of the causation of virtual objects is assuaged in a similar manner: A certain data structure, *D*, grounds a certain virtual object, *V*, so *D* determines *V*. Then, applying Yablo's argument, *V* and *D* are not causally competitive in causing their effect event, *E*, since the fact that *D* and *V* cause *E* at once does not result in overdetermination. Thus, the *Exclusion Principle* does not hold.

In sum, the causal exclusion problem does not really threaten weak virtual digitalism. Since the *Exclusion Principle* is rejected, the reasons against AC1*(W) and the conjunction of AP1*(W) and AP2*(W) do not hold. Weak virtual digitalists can still say that virtual objects have causal power without falling into the causal exclusion problem. Therefore, the arguments from causal powers (weak version) and from perception (weak version) are sound, making weak virtual digitalism an acceptable position.

2.2. Defending Virtual Digitalism from the Alternatives

After justifying weak virtual digitalism, this section compares weak virtual digitalism with its alternatives. I will consider a realist and an irrealist alternative and defend weak virtual digitalism from these views.

2.2.1. Virtual Socialism

In the realist camp, an alternative view to virtual digitalism is *virtual socialism*. According to this view, all virtual objects are real by virtue of being social objects (Ludlow, 2019: 347–351). Roughly speaking, social objects are brought into existence only if a society "takes" them to be so.

Cases include virtual organization, virtual currency, virtual tables, and virtual sushi. For instance, the users of the online game *The Sims Online* created many in-game institutions, such as The Sims Shadow Government:

> "A number of users in the Sims Online created criminal gangs, modeled on Mafia organizations for the most part, and used those gangs as instruments for the griefing (and in some cases extortion) of other players in the game. These organizations came complete with organizational charts, featuring capos and captains and various other forms of gangland titles. Some groups emerged as responses to those virtual criminal gangs, including one that called itself The Sims Shadow Government, or SSG." (*Ibid.*: 347)

As Ludlow continues, the SSG was clearly real in that the users took it to be a robust organization in *The Sim Online* society. It had a robust organizational structure with a robust communication network, even if it did not have physical offices or guns. Similarly, a virtual table is not a physical table on which we commonly place things, and virtual sushi is not edible. But they are real just because their existence is built on the function of social consensus upon which VR users agree. Hence, virtual objects are social objects and "have an existence that is quite independent of data structures" (*Ibid.*: 349).

It is also argued that virtual socialism can accommodate the putative challenges to weak virtual digitalism. This view explains the causal efficacy of virtual objects without rejecting the *Exclusion Principle*. Since the view identifies virtual objects with social objects, the causal powers of the former just are those of the latter. In one account, the causal efficacy of social objects is autonomous at the level of social explanation and irreducible to more fundamental objects (Elder-Vass, 2010; Lawson, 2013).[8] Given that social objects are causally efficacious, virtual objects are, too.[9] Also, there is no analog of the cross-play problem for virtual socialism, for virtual objects are individuated in terms of their social functions. For instance, the frisbee that Neil and Nathan play with counts as one object because it satisfies the function of what Neil and Nathan play in the game

[8] For arguments for causal autonomy, see List & Menzie (2009); Menzie & List (2010).
[9] Ludlow does not explicitly take sides in the debate of what social objects are, as they might be grounded in psychological properties, social construction, social kinds, physical properties, or some combination of the above (*Ibid.*: 364).

Frisbee, even if it is grounded in distinct data structures on their devices. Having these advantages, virtual socialism is considered a good alternative to virtual digitalism.

Although virtual socialism is appealing, I do not think this view captures the nature of virtual objects. Virtual objects can exist without any social member taking them to exist. For example, a virtual object can be generated in a decontextualized virtual world (e.g., in a virtual scenario containing only one 3D entity). The object qualifies as an existence because the underlying data structure realizes its 3D graphic. Moreover, the object need not be perceived. Its persistence condition depends on the duration of the realization. Once the data structure ceases to realize it, the object ceases to exist. No social component is involved in the object's existence. Thus, at least some virtual objects are not social. Virtual objects cannot be identified with social objects.

Admittedly, many virtual objects are socially dependent, and whether these objects exist is independent of whether they are realized by data structure. Without a society that takes SSG to be a social institution, there is no such virtual organization, even if its organizational structure and communication network in games are realized by concrete underlying data structures. However, these reasons can only characterize the nature of *social* virtual objects, rather than the nature of virtual objects. Although *some* virtual objects are social, it is not the case that *all* virtual objects are social. We can distinguish between social and non-social virtual objects and add a socially-dependent condition of existence to the former. However, whatever the condition of the social virtual objects is, it is irrelevant to the existence condition of virtual objects. Weak virtual digitalism takes social virtual objects as a subset of virtual objects, so the existence condition of each subset differs from that of the superset, namely, virtual objects. Thus, although SSG *qua* social virtual object depends on society, SSG *qua* virtual object is still grounded in data structure. Analogously, in the realm of ordinary objects, social objects fall into a subset of physical objects. A country, for instance, is a social object. It is real only if the society "takes" it to be.[10] But a country does not exist without its physical grounds, such as its people, its territory, or even its governmental infrastructure. In this sense, although a country *qua* social object depends on society, a country *qua* physical object is still grounded in its physical grounds. Just as we allow subsets of physical objects to have more grounds than their physical composition, Weak virtual digitalism is

[10] For instance, while having all the constituents of a country, Taiwan is still not globally recognized as a country because the majority of the global society does not take it to be a country.

compatible with the fact that subsets of virtual objects may have more grounds than data structures. In the case of social virtual objects, the additional grounds are social.

In sum, weak virtual digitalism is more flexible than virtual socialism. Since a virtual object may or may not be a social object, it is too strong for virtual socialism to identify virtual objects with social objects. In contrast, weak virtual digitalism can accommodate the nature of socially dependent virtual objects: A virtual object may be socially dependent if society takes it to be so. But be it social or not, a virtual object *qua* virtual object is grounded in its underlying data structure.

2.2.2. Virtual Fictionalism

Except for realism, an irrealist alternative to virtual digitalism is *virtual fictionalism*, the view that all virtual objects are non-existent fictional objects. Virtual fictionalism and virtual socialism both consider virtual objects as something we take to exist. Unlike virtual socialists, virtual fictionalists deny that our taking them to exist suffices for their being real. In McDonnell & Wildman's *virtual walt-fictionalism*, "virtual reality is a kind of walt-fiction, and our engagement with VR is not different in kind from our engagement with other forms of walt-fiction" (McDonnell & Wildman, 2019: 391). The notion of walt-fiction is based on Walton's (1990) characterization, where walt-fiction is "a representational work that has as one of its functions the role of serving as a prop in a specific game of make-believe" (McDonnell &Wildman, 2019: 391); and a prop is an item "whose existence and real features are used to guide and determine features of the imagined world of the game of make-believe" (*Ibid.*: 390). For example, two stuffed animals may serve as props that represent the parents of an imagined family. In this respect, the props are objects that really exist, but the parents are not. We simply make-believe the latter to exist when we engage with a walt-fiction.

Likewise, a virtual world is a walt-fiction. Virtual objects are the objects that "we *make-believe* to exist whenever we engage with a particular VR walt-fiction in an authorized manner" (*Ibid.*: 392). In particular, in the VR game *Longbow*:

> "[The] *images* of little grey men wearing Viking helmets, castle gates, balloons, arrows, and a bow are *real*, and […] these images serve as props in a game of make-believe. But the little grey men wearing Viking helmets aren't *real*. Rather, they only 'exist' within the

> *Longbow* game of make-believe. Once we stop playing the game, they stop 'existing'." (*Ibid.*: 392–3; original emphasis)

While dependent upon the real props, virtual objects are taken as real only within the VR walt-fiction. However, such make-believe by no means suffices for the existence of virtual objects. They are just real within the game of make-believe. But they do not have any ontological status outside of the walt-fiction.

Virtual walt-fictionalism perfectly avoids the causal exclusion problem. Since virtual objects do not really exist, their alleged causal powers are also fictional. Within the relevant games of make-believe, "these objects are taken to possess the relevant powers" (*Ibid.*: 393). Instead, the real features of the relevant props play the causal role. The fact that one has the power to shoot me within the *Longbow* game is fictional because the props underpinning one's avatar have certain powers that affect the props underpinning my avatar. Thus, virtual walt-fictionalism seems a good alternative to weak virtual digitalism.

Although virtual walt-fictionalism is a coherent view, I do not believe this view is better than weak virtual digitalism. It is quite plausible that virtual objects exist and are not reducible to the corresponding data structures because they are multiply realizable by concrete metaphysical grounds. Indeed, virtual objects may be fictional, but they need not be. Recall the make-believe game in which two stuffed animals represent the parents of an imagined family. True, the stuffed animals are not parents. We make-believe them to be parents in an imagined walt-fiction. In this sense, the "parents" are fictional objects. However, the stuffed animals are stuffed animals, the concrete objects with causal powers that we take to serve as props to represent parents and their imagined causal powers. Without the stuffed animals *qua* concrete objects, the walt-fiction cannot be realized, whereas the stuffed animals still are stuffed animals without the walt-fiction. In that sense, the "stuffed animals" are real objects.

Similarly, virtual objects may not be fictional even in a walt-fiction. Virtual walt-fictionalism claims that there are real, concrete things — e.g., images in VR sets — and concrete data structures, but there are no virtual objects — e.g., Vikings with bows. This view claims that Vikings with bows are fictional objects that we make-believe with the help of props such as the images in VR sets. But walt-fictionalism seems to misplace the fiction. Make-believing is an intentional act of imagination. We do actively engage in make-believe when playing VR games.

We imagine worlds that are not directly presented to us through VR, e.g., the Viking world of *Longbow* and its factions etc. This backstory is imagined. But apart from those imagined, there is something "given" in the simulation. For instance, the props that we use when engaging in this make-belief go beyond the pixels on the VR headset. They include virtual objects such as virtual avatars, virtual clocks, and virtual bows, which we involuntarily perceive without engaging in make-believe.[11] As a result, it is implausible to say that all virtual objects are essentially fictional objects. Virtual walt-fictionalism is true of the imagined, but not of the "given" things. It is not superior to weak virtual digitalism in explaining the nature of virtual objects.[12]

In sum, I believe weak virtual digitalism is the correct view of the nature of virtual objects. This section defended weak virtual digitalism from two alternatives, i.e., virtual socialism and virtual walt-fictionalism. The next section turns to my argument for the instantiation of virtual properties.

2.3. The Nature of Virtual Properties

I turn to virtual properties. Section 2.3.1 preliminarily defines two kinds of properties: *categorical properties* and *role properties*. Section 2.3.2 argues that many virtual properties are categorical properties instantiated by virtual objects. Section 2.3.3. argues that many virtual properties are role properties instantiated by virtual objects. In sum, many virtual properties are instantiated by virtual objects, so they are real.

2.3.1. Categorical Properties and Role Properties

[11] In their footnote 22, McDonnel & Wildman (2019) suggest that some virtual objects could represent themselves in a virtual world. In other words, it is possible that "the prop is identical to the 'virtual object'" (*Ibid.*: 392). McDonnel & Wildman take this possibility as an anomaly that is unable to save virtual realism, but I disagree. I think it leaves room for virtual realists to say that some virtual objects exist *simpliciter* and are *given* in the virtual world. For example, while playing a virtual game, "the clock in my screen's upper left is not fictional, and I do not pretend it is a real clock" (Grabarczyk, 2024: 88). The clock is given in the game, and experiencing the clock need not involve imagination or make-belief. Instead, the clock is what we involuntarily perceive. While still debatable, I think this is a good counterexample against virtual fictionalism.

[12] Admittedly, this is not the end of the realist-fictionalist debate. A virtual fictionalist might respond that realists might not be able to draw a line between the fictional and the real in a non-question-begging way. She might argue that whether a virtual object is fictional or real depends on whether our experience of it involves imagination or not, and the latter question in turn depends on what distinguishes imagination from involuntary perception. And even if we draw a line between imagination from involuntary perception, why can't the involuntary perception of the so-called virtual objects be hallucinated? To confess, I have no substantial response to this argument yet. But I think, as discussed in the previous footnote, the possibility that some virtual objects are identical to their props provides a good reason to say that not all virtual objects are fictional.

Let me characterize two kinds of properties. First, *categorical properties*. Categorical properties are typically contrasted with dispositional properties: "Dispositional properties are essentially tendencies to produce certain effects, and while categorical properties may have powers to produce effects, they are not essentially tendencies to produce them" (Pereboom, 2013: 41). As traditionally defined, for a dispositional property *F* of an object X, "there are an associated stimulus condition and manifestation such that, necessarily, X has *F* only if X would produce the manifestation if it were in the stimulus condition" (Choi & Fara, 2021: §2.1). In contrast, a categorical property does not involve conditions for manifestation. It is the quality of an object instead of its tendency.

To illustrate, solubility is a dispositional property of salt. This property may never manifest if the stimulus condition (e.g., being in the water) never occurs. On the other hand, salt has categorical properties such as *having sodium ions* and *having chloride ions*. They are not dispositions but qualities of salt.

Of course, the distinction between categorical and dispositional properties is not uncontroversial. Some argue that all properties are dispositional (Mumford, 1998, 2004; Bird, 2007), while others contend that all properties are both categorical and dispositional (Heil, 2003, 2005). Since the debate goes beyond the scope of my aim, I assume the distinction and the traditional view that all dispositional properties are dependent upon categorical properties (Lewis, 1986; Schaffer, 2005).

The distinction between categorical and dispositional properties is relevant but distinct from the distinction between essential and accidental properties. Broadly understood, "an essential property of an object is a property that it must have, while an accidental property of an object is one that it happens to have but that it could lack" (Ishii & Atkins, 2020). Specifically, whether a property is essential or accidental depends on the object that has it. *Being Socrates* is essential to Socrates but accidental to a Greek philosopher. By contrast, whether a property is categorical or dispositional depends on the nature of the property itself. Assuming the distinction, a categorical property is essentially categorical, and a dispositional property is essentially dispositional. But a categorical/dispositional property may or may not be essential to an object. At least, the distinction of categorical/dispositional properties is independent of the distinction of essential/accidental properties.

Second, *role properties*. Role properties are higher-order properties that ascribe certain broadly causal roles to properties. While role properties are properties of playing certain causal roles (or properties of having a property that plays those roles), properties to which role properties ascribe the roles are those that in fact play those roles. Call them realizer properties. Typically, we say that role properties are realized by their realizer properties.[13] For example, *being a can-opener* is the role property that ascribes the can-opening role to, say, *being stainless steel in a certain shape*. In this way, we say that *being a can-opener* is realized by *being stainless steel in a certain shape*. Also, different realizer properties can realize the same role property: *Being gold in a certain shape* can also play the can-opening role, so it can also realize *being a can-opener*. Due to the multiple realizability, i.e., the possibility that a role property is realized by distinct realizer properties, the realization relation between role properties and their realizers is not so strong as reduction or identity (Putnam, 1967; Fodor, 1974). In the literature, there are different formulations of realization. Some characterize it in terms of a functional relation (Polger, 2004), a determinable-determinate relation (Yablo, 1992), a set-subset relation (Shoemaker, 2001a, 2007, 2013), a mereological relation (Gillett, 2002, 2003, 2010), etc. Since the debate exceeds my purposes here, I stay neutral to the nature of realization and take it as a non-reductive dependence relation.

A typical way to characterize a role property rests on the Ramsey-Carnap-Lewis method, which regiments theories into Ramsey sentences (Ramsey, 1954; Carnap, 1967; Lewis, 1972). The Ramsey-Carnap-Lewis method is a method that replaces theoretical terms involved in a statement of a theory with theory-neutral terms, with the new statement being equivalent to the initial one. Starting with any theoretical term φ that occurs in a complete statement of a given theory T, one can represent this theory as saying $T(\varphi)$, where T is a complex one-place predicate. In this way, the original theory T using the term φ can be turned into the equivalent Ramsified theory $\exists x T'(x)$ that does not use the term, so that φ gets its entire meaning from the role it plays within the theory. Repeating this process for every theoretical term in T, one will get a Ramsified theory that uses no theoretical terms.[14] For example, if the original theory T involving four distinct theoretical terms says $T(\varphi, \psi, \mu, \nu)$, the Ramsified theory R will be $\exists w \exists x \exists y \exists z T^*(w, x, y, z)$. Again, these theoretical

[13] For ease of discussion, by "property *F* realizes property *G*" I mean that *F*'s instantiation realizes *G*'s instantiation.
[14] I follow Chalmers's (2018: 636–8) characterization that the non-theoretical expressions left in *R* may include observational expressions (e.g., phenomenal expressions and indexical expressions), logical and mathematical expressions, as well as certain auxiliary structural expressions (e.g., nomic and causal expressions and expressions constraining the classes of properties involved). For more discussion of what can or cannot be Ramsified, see Chalmers (2012: 409–22).

terms get their entire meaning from their roles in *R*, as *R* "encapsulates the content of the original theory *T*, so that *R* and *T* are conceptually equivalent" (Chalmers, 2018: 636).

One can apply the Ramsey-Carnap-Lewis method to characterize role properties. A role property ascribes a certain causal role to its realizers. The full predication of the role can be considered as a complete statement within folk theories, such that we can turn it into an equivalent Ramsey sentence so that the role gets its entire meaning from the full list of its predicates. For instance, *being a can-opener* ascribes the can-opening role. Presumably, we have a "can opener theory" to explain what a can opener is: If something is a can opener, it can be used to open cans. Now replace the theoretical term "can opener" with a theory-neutral term X. We get a statement that is equivalent to the "can opener theory", such that X is what can be used to open cans. So, the can-opening role gets its meaning from this Ramsified theory: What plays the can-opening role is what can be used to open cans. Therefore, *being a can opener* ascribes the can-opening role to objects that can be used to open cans. What instantiates *being a can opener* is whatever plays the can-opening role.

A more complex example may be *being a computer*, which ascribes the computer role that gets its entire meaning from the "computer theory" enlisting a full predication of computer, such as taking input data, processing the data, outputting the data, and storing the data. Replacing "computer" with a theory-neutral term Y, we get an equivalent Ramsified theory, such that Y is what takes input data, processes the data, outputs the data, and stores the data. In this way, the computer role gets its entire meaning from the Ramsified theory: What plays the computer role is what takes input data, processes the data, outputs the data, and stores the data. Therefore, *being a computer* ascribes the computer role to objects that take input data, process the data, output the data, and store the data. What instantiates *being a computer* is what plays the computer role.

Moreover, characterizing role properties through the Ramsey-Carnap-Lewis method preserves their multiple realizability. Different materials in different types of configurations can play the can-opening role and, in turn, instantiate *being a can opener*. A can opener can be made of steel or gold, or made to be one mile long or in a star shape. Similarly, a computer is realized by any format that plays the computer role. As a result, a role property can be instantiated by physically heterogeneous objects.

Some notes: First, I remain neutral to whether categorical and role properties are mutually exclusive. I am unsure if role properties should be taken as dispositional properties so that they are

mutually exclusive by definition. But to avoid complexity, I assume that none of my targets (i.e., virtual properties) is categorical and role property simultaneously. Second, I do not claim that categorical and role properties exhaust all properties. Some properties might be neither categorical nor role properties, but they are irrelevant to my argument.

2.3.2. Categorical Virtual Properties

Now, I argue for the instantiation of virtual properties. According to weak virtual digitalism, virtual objects exist. Presumably, these virtual objects seem to instantiate many properties that we ascribe to them. This sub-section argues that many virtual properties are categorical properties and are instantiated by virtual objects, and the next sub-section turns to role properties.

Chalmers (2017) points out that an implicit distinction is involved in the ordinary discourse about virtual properties. For most properties *X*, there will be a corresponding virtual property *virtual X*. Sometimes, we ascribe a predicate that denotes *X* to a *virtual X* that is literally *X*. But sometimes, we ascribe it to a *virtual X* that is not literally *X*. Chalmers (2022: 200) further calls *X* a *virtual-inclusive property* when *virtual X* is literally *X*, and *virtual-exclusive property* otherwise. To illustrate, we ascribe "being a kitten" to a virtual object while agreeing that it is not actually a kitten because a kitten has to be a DNA-based biological system, whereas a virtual kitten need not be. In this sense, we attribute a virtual property that is different from *being a kitten* to the virtual object. "Being a chair" and "being a car" are predicates of this sort, and the corresponding properties *being a chair*, *being a car*, as well as *being a kitten* are virtual-exclusive properties. On the other hand, we ascribe "philosopher" to a virtual object while agreeing that it is actually a philosopher, since whether an object is a philosopher depends "only on the abstract causal organization and the mental properties of a situation" (Chalmers, 2017: 325). In that sense, we employ the predicate to attribute to the virtual object the property that is identical to the property had by physical objects in the same abstract causal organization or situation of the mental. "Being a library" and "being a club" are predicates of this sort, and the corresponding properties *being a library*, *being a club*, as well as *being a philosopher* are virtual-inclusive properties.

Chalmers's characterization of virtual-exclusive properties overlaps my characterization of categorical properties. Usually, a virtual exclusive property is a categorical property. A categorical property is wholly qualitative in its object. *Being a kitten*, for instance, requires a quality that is categorically defined, and the quality in question is a biological or physical quality. By definition,

a biological or physical quality is not a virtual quality, so it excludes any virtual counterpart that could be identified as *being a kitten*. However, this does not follow that we cannot correctly call anything a "kitten" through ordinary language while using VR. "Kitten" does predicate a certain property of virtual objects, but the property is distinct from *being a kitten*. Call this property *being a virtual kitten*.[15] *Being a virtual kitten* is also wholly qualitative in its object. The quality is neither biological nor physical. Instead, it is categorically unique to virtual objects: Only virtual objects have such a quality. In short, through ordinary language, we use the predicate "being a kitten" to ascribe distinct categorical properties to virtual and non-virtual objects. This explains why sometimes a *virtual X* is not *X*. For a detailed distinction, I call the virtual-exclusive *X* the *categorical physical properties* and the corresponding *virtual X* the *categorical virtual properties*.

Plausibly, virtual objects instantiate categorical virtual properties. First, it is widely accepted that ordinary objects instantiate categorical physical properties. An ordinary object would not be the same object if it did not instantiate the categorical property it actually instantiates. A kitten is a kitten because it instantiates the categorical physical property *being a kitten*. Second, if categorical physical properties are instantiated by ordinary objects, then reasonably, categorical virtual properties are instantiated by virtual objects. This conditional appeals to philosophers' motivation for property attribution. By default, properties are entities attributed to objects, helping us categorize their natures. Attributing categorical physical properties to ordinary objects helps us explain why numerically distinct objects have the same or different natures. For example, we attribute the same categorical physical property, *being a kitten*, to black kittens and white kittens in order to explain the fact that they have a common biological quality, e.g., a certain DNA pattern that determines the biological category of *Felis catus*. The same reason applies to categorical virtual properties: We attribute the same categorical virtual property, *being a virtual kitten*, to black virtual kittens and white virtual kittens in order to explain the fact that they have a common virtual quality, e.g., certain arrays of bits that determine the category of virtual objects that we call virtual kittens. Given these reasons, it is plausible that virtual objects do instantiate categorical virtual properties.

[15] I assume that both *being a kitten* and *being a virtual kitten* are categorical properties, though the assumption is not uncontroversial. Also, *being a virtual kitten* should be considered as a *single* property. It should not be taken as a complex property that combines the property of *being a kitten* and the property of *being virtual*.

2.3.3. Virtual Properties as Role Properties Instantiated by Virtual Objects

Except for categorical virtual properties, many virtual properties are role properties instantiated by virtual objects.

Chalmers's characterization of virtual-inclusive properties overlaps my characterization of role properties. To reiterate, a virtual-inclusive property depends "only on the abstract causal organization and the mental properties of a situation" (Chalmers, 2017: 325). A virtual calculator, for instance, is a calculator because *being a calculator* depends only on the relevant organization, which is also present in virtual worlds. Also, a virtual philosopher is a philosopher because *being a philosopher* depends only on how we take a philosopher to be. These properties do not depend on the categoricity of their objects. Instead, they depend on the roles their objects realize. A calculator is one that plays the calculator role, and the role can be Ramsified to get its entire meaning from the full list of its predicates. For instance, a calculator role is the role of performing arithmetic operations on numbers, including addition, subtraction, multiplication, and division. Whether an object is made of metal, plastic, or even embedded in an advanced computer, it *is* a calculator once it plays the calculator role. Also, a philosopher is one that plays the philosopher role, which can be Ramsified as being professionally interested in philosophical ideas and trying to do something with these ideas. Similarly, not only can a human be a philosopher, but a tortoise or a robot can also be a philosopher if it plays the philosopher role.[16]

Virtual objects instantiate many role properties. It is plausible that many role properties are multiply realized by categorical physical properties and categorical virtual properties, as their Ramsey sentence is neutral to which categorical properties play the role. In addition, realization is metaphysical necessitation: The instantiation of realizer properties necessitates the instantiation of the associated role properties. If the realizer properties are instantiated, the role properties are also instantiated. Moreover, as virtual objects do instantiate categorical virtual properties, it is plausible that some role properties realized by categorical physical properties are also instantiated in terms of the instantiation of the associated categorical virtual properties. A physical calculator instantiates the role property *being a calculator* because the instantiation of its categorical physical bases necessitates the instantiation of the role property. The same role property can be instantiated by a virtual calculator in that the instantiation of its categorical virtual bases necessitates the

[16] Imagine that in reality there is a tortoise behaving like Master Oogway featured in *Kung Fu Panda*, or an intelligent robot behaving like human philosophers do. One might be willing to call them philosophers.

instantiation of *being a calculator*. Similar analyses apply to many role properties, including *being a philosopher*, *being a can-opener*, *being a computer*, etc. In this respect, virtual objects do instantiate many role properties that are also instantiated by physical objects.

In sum, many virtual properties are categorical virtual properties or role properties, and the two kinds of properties help explain why some properties are virtual-exclusive but some are virtual-inclusive. Also, we have good reasons to believe that virtual objects instantiate these virtual properties.

2.4. Conclusion

I conclude that virtual objects are real and instantiate virtual properties. Sections 2.1 and 2.2 argued that virtual objects are real by defending weak virtual digitalism: Virtual objects are real and are metaphysically grounded on data structures. Section 2.3 argued for the instantiation of virtual properties by explaining the nature of different kinds of virtual properties. Many virtual properties are categorical virtual properties or role properties and are instantiated by virtual objects.

This conclusion is set as a preliminary of my arguments for virtual veridicalism in the following chapters. To perceive an object requires the object to exist. So far, my argument in this chapter suggests that there are virtual objects and properties that we can perceive in a virtual world. Presumably, some virtual objects and properties (if not all) are obviously perceived by VR users, otherwise we would not be talking about them. However, even if we do perceive them, it is unclear whether we perceive them as they really are. For instance, this chapter suggests that the virtual avatars, the virtual clocks, and the virtual bowls are real in the *Longbow* game because players can involuntarily perceive them. Yet, the fact that players do perceive these virtual objects does not imply that these players perceive them *as* having the properties that they do have, e.g., having a certain color or being in a certain size or shape. Having answered *whether* we can perceive virtual objects and properties in this chapter, we still need to look more closely at *how* we perceive them. In the following chapters, I will argue that we perceive virtual objects *veridically*, i.e., we perceive them *as having the properties that they do have*.

Chapter 3: On the Veridicality of VR Experiences

In Chapter 2, I argued that virtual objects exist and instantiate virtual properties. In a virtual world, we probably interact with these objects and properties. We could see a virtual apple as red. We could use a virtual weapon to win an e-sport game. We could also feel as if a virtual avatar is our own body. The "apple," the "weapon," and the "avatar" exist in the virtual world. As I also suggested, many properties in virtual worlds are different in kind from their non-virtual counterparts. The virtual "apple" is not an apple, as it is not edible in the normal sense of "edible". The virtual "weapon" may not be a weapon, as it cannot hurt others. By contrast, the virtual "avatar" may be a real avatar if it plays the avatar role. As I argued in Chapter 2, these differences are accounted for by the metaphysics of categorical properties and role properties.

However, the metaphysical account of virtual objects and properties I offered in Chapter 2 does not imply that we perceive virtual objects *as they really are*. I have suggested that we perceive virtual objects in VR, which requires that such objects exist but not that we perceive them as they really are. This raises a question that is left from Chapter 2: What do we experience while perceiving virtual objects? For example, while perceiving a virtual apple that is supposedly inedible, we might perceive it either as inedible or as edible; or while perceiving a virtual weapon that is supposedly innocuous, we might still perceive it either as innocuous or as being powerful enough to harm others. Presumably, there is a significant difference in these pairs of experiences because only the first experience of each pair (i.e., the inedible experience and the innocuous experience) conveys the information about the object *as it really is*. Hence, regarding perception in VR, several questions arise: Do our perceptual experiences about virtual objects and properties tell us anything true about the virtual world? Do our perceptual experiences in a virtual world (VR experiences, for short) represent virtual objects and properties *as they really are*? More specifically, are our VR experiences veridical representations of these entities or merely a grand illusion? These questions are central to my defense of *virtual veridicalism* in this dissertation, the view that many VR experiences are veridical rather than illusory. [1]

This chapter aims to answer these questions by offering my first argument for virtual veridicalism. First of all, some terminology. An *experience* is a mental state that has a phenomenal

[1] The distinction between illusion and hallucination is irrelevant to the topic. By *illusion* I mean falsidical experience.

character. The *phenomenal character* of an experience is what it is like for the subject to have that experience. *Perceptual experiences* are restricted to experiences in sense modalities. *VR experiences*, as I define them, are perceptual experiences in a virtual environment simulated by VR technology, namely, a virtual world. They are perceptual experiences of virtual objects and properties.

The chapter proceeds as follows. Section 3.1 introduces virtual veridicalism and the relevant debate. Sections 3.2 – 3.4 argue for this view by a series of thought experiments from perceptual variation without misperception. Section 3.2 argues for perceptual variation without misperception in color perception; Section 3.3 argues for perceptual variation without misperception in size and shape perception; Section 3.4 further argues for perceptual variation without misperception in color, size, and shape perception across virtual and non-virtual environments; Section 3.5 addresses the scope of virtual veridicalism, suggesting that it should encompass not only perfect virtual worlds (i.e., virtual worlds that *completely* screen off non-virtual things and then appear the same as non-VR environments) but also mixed worlds (i.e., virtual worlds that *partly* screen off non-virtual things and then appear the same as non-VR environments). While this chapter focuses on visual experiences, I believe the conclusion generalizes to experiences in other modalities.

3.1. Virtual Veridicalism vs. Virtual Illusionism

Due to advanced simulation technology, many imaginative scenarios are realized in VR. A VR user, Alice, could feel as if she is located in Wonderland, manipulating an avatar to venture into a queen's croquet ground. She could experience the avatar as her own body, given appropriate stimulation. She might also feel she could escape from the croquet ground using the avatar. However, when Alice is asked whether her experiences in the simulation were veridical, she would probably answer without doubt that they are all illusory. Part of the reason would be the intuition that what she experienced does not match reality: She experienced a croquet ground, but no land was there; she experienced the avatar, but it was not her body; she thought she escaped, but actually she did not move, etc.

Many researchers claim that VR experiences are illusory in these ways. They characterize the illusions that a user would experience in the Wonderland scenario as follows:

- Place illusion: "the strong illusion of being in a place in spite of the sure knowledge that you are not there." (Slater, 2009: 3551)
- Plausibility illusion: "the illusion that what is apparently happening is really happening (even though you know for sure that it is not)." (Slater, 2009: 3553)
- Body ownership illusion: the illusion of body ownership with respect to the avatar in virtual reality. (Maselli & Slater, 2013)
- Power illusion: the illusion that one feels as if one's virtual gesture is powerful to create agency in virtual reality. (Langer, 1953: 175)

These researchers would probably suggest that Alice suffered from these illusions: She illusorily experienced herself as being located in Wonderland, which she knew does not exist; she illusorily experienced something happening at the croquet ground; she illusorily experienced the avatar as her own body; she illusorily experienced that her effort using the avatar would help her escape. The mismatch between what Alice perceived and what was happening *out there* in her environment results in the illusoriness of her VR experiences.

Generalizing this list, researchers argue for *virtual illusionism*, claiming that *all* VR experiences are illusory (Slater, 2009; Maselli & Slater, 2013). In one version, the illusoriness of VR experiences consists in the fact that VR users perceive virtual things as having properties that no relevant object has:

> "In VR the perception arises from the interpretation of arrays of illuminated pixels as a flower located in extended space, a flower that the perceiver can walk around and admire from multiple viewpoints. *The flower in reality is a physically existing object*. From the moment that the light from the illuminated pixels in the case of the virtual flower reaches the eyes of the beholder, the perceptual process can be considered as the same as when the light from the real flower reaches the eyes. From this moment we can consider that the brain takes over, and can activate in the same way independently of the source of the light reaching the eyes. We refer to the perception of the virtual flower as illusory, not because the perception itself is illusory, but due to the source of the light. *The perception itself is real, but the perceived flower is not real, from the standpoint of an observer outside of the VR*." (Slater & Sanchez-Vives, 2022: 4; italics mine)

The object perceived by VR users, though looking like a flower, is not a (physical) flower, since it does not satisfy the physical criterion of being a flower. It is an illusion if the user perceives the object as a flower, as the object does not possess the property of *being a flower*. Furthermore, as Slater and Sanchez-Vives suggest, since flowers are physical, no virtual object could possess the property of *being a flower*. This implies that, if a VR user perceives any virtual object as a flower, she must suffer from an illusion.

Virtual illusionism is intuitive if we think that VR users perceive virtual objects as having physical properties that no virtual object really has. However, this is not necessarily the case. It is debatable that no virtual properties are genuinely perceived in VR experiences. Casting doubt on virtual illusionism, Chalmers contends that what users perceive in VR may really be *out there*:

> "While it's true that VR can involve an illusion, it doesn't have to, and for many users it won't involve an illusion. The users' perception of place, of plausibility, of power, and of embodiment needn't be illusory. It will often be an accurate guide to their virtual world. In many cases, users have the sense of being in a virtual (not a physical) place, and *they really are in that virtual place*. Users may have the sense that things are happening in a virtual (not a physical) world, and *those things really are happening in the virtual world*. Users may have the sense of having a virtual body (not their own physical body), and *they really do have that virtual body*. They have the sense of performing virtual actions that they really perform. None of these things need be illusions." (Chalmers, 2022: 205; italics mine)

This passage suggests that VR users do not necessarily perceive virtual objects as having physical properties. Instead, they can perceive them as having virtual properties. Recall the Wonderland scenario. Alice may experience herself as being located in a *virtual* Wonderland; she may experience something happening at the croquet ground as a *virtual* event; she may experience the avatar as her *virtual* body; she may also experience that her effort using the avatar would help her virtual body escape from the dangerous *virtual* environment.

Furthermore, as supported by the conclusion of Chapter 2, what Alice perceives is really happening *out there*. Wonderland is a virtual world that contains many virtual objects and properties. The queen, the croquet ground, Alice's virtual body, and even Wonderland itself are

virtual objects realized by underlying data structures. Their characteristics and mutual relationships are virtual properties. Given these virtual entities, it is quite reasonable to say that *there is a virtual croquet ground*, that *those things are happening in the virtual croquet ground*, that *Alice has that virtual body*, that *Alice can escape using her virtual body*, etc. For these reasons, it is no longer plausible to say all of Alice's experiences in Wonderland are illusory. At least some of them can be veridical. This view is *virtual veridicalism*.

To summarize, virtual illusionism and veridicalism disagree on whether VR experiences can be veridical:

Virtual illusionism: All VR experiences are illusory.
Virtual veridicalism: Many VR experiences are veridical.

This chapter defends virtual veridicalism in terms of what I call *lenient veridicality*, as opposed to *strict veridicality*. Let me elaborate on the difference.

Experiences have *contents*. The content of an experience is its veridicality condition, i.e., the condition under which the experience is veridical. On the face of it, there are many ways to associate content with an experience. There is a kind of content we associate with experiences that represents the property that appears in the experiences' phenomenology. In the case of reddish experience, such a property is phenomenal redness (see also Chalmers, 2006). Call this kind of content the *strict content*. In addition, there is another kind of content we associate with experiences that represents the property that is partly determined by the appropriate contingent relationship between the experiences and the external world. Call this latter kind the *lenient content*. For example, some might argue that *having a surface reflectance profile that reflects light with a wavelength of 700nm* (C_{700}, for short) is the lenient content of reddish experience. Due to the distinction between strict and lenient contents, there are also two senses with respect to which an experience is veridical. I call the veridicality condition with respect to strict content *strict veridicality* and the veridicality condition with respect to lenient content *lenient veridicality*. An experience is *strictly veridical* iff its strict content is true, while it is *leniently veridical* iff its lenient content is true.

An experience can be strictly falsidical but leniently veridical. When we see a ripe tomato, we typically perceive it as red. Apparently, the strict content of that tomato experience represents

phenomenal redness, so it is true iff the ripe tomato instantiates phenomenal redness. However, as argued by many philosophers (e.g., Boghossian & Velleman, 1989; Chalmers, 2006; Mendelovici, 2013; Cutter, 2021), it is implausible that the ripe tomato does instantiate phenomenal redness. That is, the tomato experience is strictly falsidical.

Rather, it is undeniable that the ripe tomato instantiates some property that bears an appropriate contingent relationship to our reddish experience so that we *typically* perceive the tomato as red. As defined above, such a property is lenient content, so the reddish experience is veridical iff the tomato instantiates such a property. This chapter will not specify the nature of lenient content, which I will spend the entire Chapters 4 and 5 to accomplish. In addition, I leave open the list of the constraints of the search for the "appropriate contingent relationship" of the lenient content. I believe the list largely depends on how we take our experience to be useful in practice. But I also believe the list should contain these three constraints: The appropriate contingent relationship must 1) help our successful interaction with the environment, 2) make relevant utterances intelligible, and 3) justify common knowledge.

To illustrate, consider a common case where drivers check the traffic through the rear mirror. Although the cars in the mirror appear in the front, an experienced driver usually sees them as behind her car. Commonsensically, we do not consider the driver's experience illusory, since it satisfies the following constraints. First, the driver navigates the environment successfully through her mirror experience. If what the experience conveys to the driver is not reliable with respect to her environment, she will probably not be able to react appropriately. A car accident might occur if she perceived the car in the mirror as in the front while reverse parking. Second, the mirror experience makes our conversation intelligible. Presumably, we employ linguistic apparatus to describe the environment in the way we perceive it. The mirror experience serves as a common ground for conversation. Since drivers and passengers perceive the cars in the mirror as behind them, it makes sense to say "These cars are behind us" based on their mirror experience. It would also be odd if a passenger said "Watch out for the cars in front of us!" as the sentence is unintelligible in the context. Third, the mirror experience plays a crucial role in justifying our knowledge. As many believe that perceptual experience provides an immediate justification for knowledge (Pryor, 2000; Huemer, 2001; Kriegel, 2021), our experience is best explained as veridical if the justification is undefeated. In particular, assuming that the driver is aware of the mirror's presence, her knowledge that a car is behind her is immediately justified by her mirror

experience whose content involves something like <a car is behind me>. Without a defeater against the knowledge, it is plausible that the mirror experience conveys the truth to the driver. In sum, given the satisfaction of these conditions for practical purposes, it makes sense to say that the driver's mirror experience is leniently veridical, even if it is not strictly veridical.[2] In this chapter, I mean lenient veridicality by "veridicality" unless specified otherwise.

I clarify other nuances before proceeding. First, my defense centers on *low-level properties*, including colors, locations, and shapes. I also assume that for any low-level property, while being perceptually experienced, there is a phenomenal character in the experience representing that property. Second, I leave open what kinds of contents are represented in perceptual experiences (e.g., singular vs. general contents) and what is the correct metaphysics of content (e.g., Russellian vs. Fregean contents). This chapter only requires that experiences have lenient content, whatever its nature is. Staying neutral on the issues of content ensures the compatibility of my argument and a variety of positions in the philosophy of perception.

This chapter specifically defends *virtual veridicalism about low-level properties by lenient veridicality,* i.e., many VR experiences of low-level properties are leniently veridical. For ease of terminology, in the rest of this dissertation, I mean this particular version simply by "virtual veridicalism" unless specified otherwise. My argument exploits a series of thought experiments from perceptual variation without misperception,[3] the phenomenon in which two groups of perceptual experiences of the same type veridically represents different properties in different circumstances. I argue that perceptual variation without misperception is possible in low-level perception, such as color, size, and shape, and that the VR experience is an instance of that possibility. Here is an overview of my main argument:

PV1 Perceptual variation without misperception is possible.

[2] My distinction between the strict and the lenient criterion of veridicality is similar to Chalmers's (2006) between perfect and imperfect veridicality. In fact, my strict criterion coincides with his perfect veridicality, requiring a perfect match between the properties instantiated in experience and the properties instantiated by the perceived object. However, while Chalmers's imperfect veridicality requires a particular view of perceptual content (viz., Fregean content), my lenient criterion does not assume any particular theory of content so is more tolerant to views in the relevant debate.

[3] By *misperception,* I mean falsidical perception.

PV2 If perceptual variation without misperception is possible, then the perceptual variation across virtual scenarios and ordinary environments is an instance of perceptual variation without misperception.

PV3 If the perceptual variation across virtual scenarios and ordinary environments is an instance of perceptual variation without misperception, then virtual veridicalism is true.

PV4 Therefore, virtual veridicalism is true.

Sections 3.2 and 3.3 justify PV1: Section 3.2 concerns color perception, and Section 3.3 addresses size and shape perception. Section 3.4 defends PV2. PV3 follows from the definition of virtual veridicalism: If some experience type represents virtual properties without misperception, many VR experiences are veridical. Section 3.5 generalizes the scope of virtual veridicalism to scenarios of mixed reality. While this chapter focuses on visual experiences, I believe the conclusion generalizes to experiences in other modalities.

3.2. Perceptual Variation without Misperception in Color Perception

Perceptual variation is the phenomenon in which two groups of phenomenally identical perceptual experiences are reliably generated by different properties in different circumstances, and perceptual variation *without misperception* adds that neither group of experiences is falsidical. The possibility of perceptual variation without misperception has long been debated in the philosophical literature on color perception. The famous thought experiment of spectrum inversion (Locke, 1689/1975) puts forward the idea that it is possible for different subjects to veridically perceive objects that reflect lights with distinct wavelengths as having the same color. Particularly, it is possible that a subject perceives a ripe tomato as uniquely red while her phenomenal duplicate perceives grass as uniquely red, and neither suffers from misperception.

Refined versions of spectrum inversion are ubiquitous in contemporary philosophy of mind. I adapt the popular Inverted Earth (Block, 1990, 1996) for my argument. Suppose Sally is a normal human perceiver on our planet, Earth. While seeing a basketball, Sally perceives it as being orange in color. Now, imagine a far-off planet on which there is an individual, Cally, who is a phenomenal duplicate of Sally, so Cally's basketball experience also involves an orangish character. The planet

Cally inhabits is physically similar to Earth in most respects, yet it is also fundamentally different from Earth in the following ways:

> *Color-inverted Earth.* On Color-inverted Earth, grass has the surface reflectance profiles that cause Sally's reddish experiences; strawberries have the surface reflectance profiles that cause Sally's greenish experiences; bananas have the surface reflectance profiles that cause Sally's bluish experiences, etc. Nonetheless, Color-inverted Earthlings, including Cally, perceive things in such a way that their grass experience involves the greenish character, their strawberry experience involves the reddish character, and their banana experience involves the yellowish character. Likewise, although basketballs on Color-inverted Earth have the surface reflectance profiles that cause Sally's turquoise-bluish experiences, Cally's basketball experience involves the orangish character. (Case adapted from Block, 1990)

If the case is possible, the difference between Sally's and Cally's color experiences is an instance of perceptual variation in color perception, as their phenomenally identical orangish experiences are reliably generated by distinct types of physical properties. Now, I am going to argue that:

CE1 *Color-inverted Earth* is possible, and

CE2 Neither Sally nor Cally suffers from misperception in color.

CE1 is considered true by many philosophers, as the possibility of color inversion does not violate our intuitions about the external world (Locke, 1689/1975; Peacocke, 1983, 1992; Shoemaker, 1982, 1996). The way Cally interacts with her surroundings is not counterintuitive. It is feasible for distinct properties to reliably produce a specific type of experience in various ways. Specifically, the color experiences can be generated in the Earth's way and the Color-inverted Earth's way. No current scientific data suggest that Cally's color perception is impossible. For this reason, CE1 is probably true.

A straightforward objection to CE1 is to deny the possibility of color inversion. As Byrne & Hilbert argue, the possibility renders the following statements inconsistent:

BH1 The phenomenology of Cally's color experiences is inverted with respect to Sally's.

BH2 Sally's color experiences are veridical.

BH3 Cally's color experiences are veridical.

BH4 For all possible subjects S_1, S_2 and all possible worlds w_1, w_2, if S_1 is having a visual experience in w_1 and S_2 is having a visual experience in w_2, then these experiences represent the same (physical) color properties iff they are the same in color phenomenology.[4] (argument adapted from Byrne & Hilbert, 1997: 267–272)

BH1 stands for the possibility of *Color-inverted Earth*. If possible, then Sally and Cally have color experiences with the same phenomenology while perceiving things that generate color experiences with inverted phenomenology in a single person. BH2 and BH3 will be defended later (CE2). I assume them for the present purpose. As for BH4, introspection shows that visual phenomenology does not sever the phenomenal character from what it represents, so the sameness in one suffices for the sameness in the other. BH1 – BH4 together entail that Sally's and Cally's phenomenally identical experiences represent the same property on their respective planets. This consequence contradicts the setting of *Color-inverted Earth*, according to which Sally's and Cally's phenomenally identical experiences are produced by distinct physical properties. Therefore, one of the premises in Byrne & Hilbert's argument should be rejected.

To resolve the inconsistency, Byrne & Hilbert reject BH1: If BH1 is true, we could not describe Sally's and Cally's phenomenally identical experiences as differing in some elements of representation that ascribe different represented properties. This is unacceptable for Byrne & Hilbert, as proponents of *Color-inverted Earth*, i.e., those who believe that there are orange-feeling experiences that are not orange-representing, "owe us an account of what is common to all and only [orange]-feeling experiences" (*Ibid*.: 283, fn. 16).

However, BH1 can be saved by a sufficient account of the commonality of a certain feeling. For example, Peacocke (1983) proposes that the commonality lies in the common *sensational*

[4] BH4 is slightly revised from Byrne & Hilbert's (1997) original text. The original version is the Necessity Principle: For all possible subjects S_1, S_2 and all possible worlds w_1, w_2, if S_1 is having a visual experience in w_1 and S_2 is having a visual experience in w_2, then these experiences are the same in *color content* iff they are the same in color phenomenology. Here are my reasons for not presenting BH4 in this way. First, since Byrne & Hilbert also assume color physicalism and that "two experiences are the same in color content just in case they represent the same color properties instantiated at the same (viewer-centered) locations" (*Ibid*.: 267), BH4 does not alter the original meaning of the Necessity Principle. Second, I intend to avoid discussing contents in this chapter, as they affect what perceptual veridicality is. For the present purpose, it suffices to present BH4 in the way I do.

property instantiated in regions of the visual field in that type, so Sally's and Cally's orangish experiences share the same sensational property even if they represent different (physical) color properties. Shoemaker (1996, 2006) also contends that the commonality of orangish experiences is explained by the fact that they represent objects as having a common *phenomenal property*, i.e., *phenomenal orange-ness*. These accounts separate the analysis of phenomenal characters from the physical nature of stimuli. Be they sufficient, both accounts save BH1 without falling short of the problem pointed out by Byrne & Hilbert.

Moreover, Byrne & Hilbert's reason for BH4 is based on a sort of tracking theory of intentionality, which is roughly the view that what an experience represents just is what it tracks. However, as mismatches between perceptual representation and tracking relation pervade in a variety of dimensions (Hardin, 1988; Akins, 1996; Pautz, 2006, 2014; Mendelovici, 2018: Ch. 3), we have good reasons to reject the tracking theory in this manner.[5,6] For these reasons, I suggest rejecting BH4 to resolve the inconsistency while preserving BH1. Thus, *Color-inverted Earth* is indeed possible. CE1 holds.

I turn to CE2, arguing that perceptual variation on Color-inverted Earth does not suffer from misperception. Although Byrne & Hilbert suggest a wrong solution to the inconsistency, I agree with their acceptance of BH2 and BH3. BH2 and BH3 constitute CE2, as CE2 states that Sally and Cally perceive color veridically. I defend CE2 by justifying BH2 and BH3.

Given CE1, orangish experience is reliably generated by distinct physical properties across Earth and Color-inverted Earth. One might ask: Is Sally or Cally perceiving veridically? There are three options. First, neither-ism: Neither Sally nor Cally perceives color veridically. Second, one-

[5] For instance, Pautz (2006, 2014), Hardin (1988), and Akins (1996) argue that there is a "structural" mismatch between perceptual representation and the tracking relation. Specifically, there is a mismatch in the similarity and difference between what a set of experiences represent and the similarity and difference between the properties that they track. Mendelovici (2018) argues that identifying perceptual representation with the tracking relation both wrongly predicts unwanted materials in the representation and wrongly omits essential materials from the representation.

[6] Certainly, tracking theorists have proposed possible solutions to these objections. They might revise the theory in that color experiences do not represent single physical color properties but a *disjunction* of physical color properties (e.g., Byrne & Hilbert, 2003). They might also argue that what perceptual experience tracks is not first-order properties but second-order properties (e.g., Cutter, 2021). But none of the solutions saves BH4 from the objections. Placing all the physical properties that orangish experience tracks into a disjunction makes the first solution *ad hoc* to account for "what is common to all and only [orange]-feeling experiences" (*Ibid*.). The second solution, however, changes what Byrne & Hilbert initially means by "color properties" from physically first-order properties to non-physical, second-order properties.

ism: Only Sally perceives color veridically.[7] Third, both-ism: Both perceive color veridically. The argument for CE2 consists of two stages. The first stage argues for BH2, where I reject neither-ism. The second stage argues that if BH2 is true, then BH3 is true, and then one-ism is rejected. Thus, both BH2 and BH3 are true, i.e., both-ism and then CE2 is true.

Stage 1. Presumably, Sally's color perception is considered systematically veridical by the lenient criterion of veridicality because it is good enough for practical purposes. In particular, the fact that Sally's experience represents the basketball as orange enables her to detect the ball's motion more quickly, makes her ball discourse intelligible, and grounds her knowledge about it. Because of the satisfaction of these conditions, the multivariate interactions between Sally and her environment are explained in that her color experience is veridical. Since Sally interacts well with her environment, BH2 is plausibly true.

Neither-ism must deny BH2. They would embrace an error theory of color perception, according to which Sally's color perception suffers from systematic misrepresentation, arguing that the error theory can equally accommodate the reasons for BH2 above. Boghossian & Velleman (1989) argue that our color discourse remains intelligible despite systematic misperception. Presumably, we take an ordinary statement "The sun rises" to be true due to our visual experience while knowing that the statement is never true because the sun does not literally "rise" from the horizon. Analogously, there is a systematic error in what we perceive with respect to what it really is. As they continue, the systematic error does not prevent our ordinary discourse from using color concepts to ascribe color properties to objects, since the commonality of our color experience still enhances the informativeness of color discourse. We can make sense of the recognition or categorization of objects by ascribing color to them, even if those objects do not possess such color properties.

In addition, Mendelovici (2013) contends that reliability, rather than veridicality, suffices for perceivers' successful interaction with the environment. Specifically, reliable misperception helps perceivers survive and flourish by enabling their capacities of object discrimination and re-identification through time. For example, object discrimination aids one in distinguishing the

[7] Logically, there are two one-ist positions on the table: One is that only Sally's color perception is veridical, while the other is that only Cally's is. For the present purpose, it does not matter who exactly perceives veridically. The argument is to reject that one is privileged in perceptual veridically over the other. So, this stage assumes one-ism to be presumably the position that only Sally's color perception is veridical, while Cally suffers from misperception in color.

objects that humans have evolved to avert from those we tend to approach, and object re-identification assists one in targeting the same object on different occasions without costing extra cognitive resources. These capacities can be built even if our perception suffers from systematic errors, suggesting that perceivers can successfully interact with the environment without their perception being veridical. If this line of argument succeeds, the reasons for BH2 are undermined, as the veridicality of perception is no longer needed to explain the advantages of appropriate subject-world relationships (see also Chalmers, 2006; Cutter, 2021).

However, these arguments are made for neither-ism with respect to strict veridicality but not with respect to lenient veridicality. The neither-ists consider color experience systematically erroneous because external objects "do not have the properties our color experiences represent them as having" (Mendelovici, 2013: 421). Put differently, Sally's color experiences are systematically illusory because external objects never possess the phenomenal color properties that appear in her experiences. This is precisely strict veridicality: An experience is strictly veridical iff its object instantiates the property that appears in the experience's phenomenology. Strict veridicality requires more than lenient veridicality does. People might not be sophisticated enough to distinguish phenomenal color properties from external color properties and take the standard for veridicality as high as neither-ists do. Analogously, the conflict between the scientific and intuitive truth-values of "The sun rises" does not necessarily force us to conclude that our visual experience is false. Indeed. The statement is not strictly true because the sun does not really "arise" from the horizon. But since the statement still yields successful interaction with the environment, intelligible description, and justifies common knowledge, we can still say it is leniently true. The conflict between the scientific and intuitive truth-values may be dissolved once the standard of truth is clarified. Since my argument targets lenient veridicality, BH2 is not threatened by the neither-ist objections, even if Sally's color experience is not strictly veridical.

Stage 2. If BH2 is true, then BH3 is true, and one-ism is false. Sally's color experience satisfies the constraints of lenient veridicality, so it is veridical. Specifically, the experience yields Sally's successful interaction with her environment, makes sense of the color talk, and justifies Sally's knowledge about the objects. Since Sally's color experience is leniently veridical, Cally's experience is just as well positioned to count as leniently veridical as Sally's experience.

Cally's color experience does satisfy lenient veridicality. First, the experience does not hinder her from appropriately interacting with her environment. Perceiving the basketball as

having a specific color helps orient Cally's action to the ball and enriches her ability to interact with it. Cally could chase and catch the ball more efficiently using her color experience. She could also detect the ball's motion easily if the basketball reflects light with a wavelength different from that reflected by the background. These accomplishments are not impeded even if Cally's color experience is inverted relative to Sally's. Second, Cally's color experience also makes sense of the color talk in her community (viz., Color-inverted Earthlings). Since Cally and her community reliably perceive the basketball as orange, the talk that the basketball is orange is intelligible for Color-inverted Earthlings. Their term "orange" might refer to a specific surface reflectance profile that reliably produces turquoise-bluish experience in Earthlings or something else. Whatever its reference, the term used by Color-inverted Earthlings has a commonly accepted meaning. In this way, the intelligibility of Color-inverted Earthlings' color talk is not hampered by color inversion. Third, Cally's color experience provides reliable grounds for knowledge about the color of objects. When Cally states that the basketball is orange, she expresses something about her environment. Such a statement is agreed upon as knowledge by her community because its reliable linkage with the experience suffices for justifying the statement, as most of the members would state the same while having the same experience. Hence, having color-inverted experience does not prevent Cally and Color-inverted Earthlings from gaining knowledge about colors. Due to the satisfaction of lenient veridicality, Cally's color experience is leniently veridical. BH3 is true.

Given BH2, one-ism must deny BH3, otherwise the view is false. For a one-ist, there should be a difference in Sally's and Cally's color experiences so that the former's veridicality is privileged over the latter's. If the one-ist finds it out, BH3 can be rejected even if BH2 remains true. But what is a good candidate that not only makes Sally's experience satisfy lenient veridicality but also makes Cally's experience fail to satisfy it? As far as I can imagine, there is no appropriate candidate for one-ism. An appeal to strict veridicality does not save one-ism, since it even denies BH2. Without categorial difference in the veridicality between Sally's and Cally's color experiences, it is arbitrary to privilege one's veridicality over the other's. Therefore, it is plausible that if BH2 is true, then BH3 is true.

Since neither-ism and one-ism are rejected, both-ism is true, i.e., Sally's and Cally's color experiences are both veridical. Since BH2 is true and if BH2 is true then BH3 is true, both BH2 and BH3 are true, i.e., CE2 is true: Perceptual variation in color across Earth and Color-inverted Earth does not involve misperception. CE1 and CE2 constitute a restricted version of PV1: *In color*

perception, perceptual variation without misperception is possible. The following section turns to size and shape perception.

3.3. Perceptual Variation without Misperception in Size and Shape Perception

Let's continue with Sally, a normal human perceiver on Earth. While seeing a basketball, Sally perceives it not only as being orange. She also perceives it as being of a specific size (called α-size) and being round. Now, imagine another two far-off planets. On each planet, there is Sally's phenomenal duplicate. They are Dally and Gally, respectively. Sally shares all the phenomenal characters with them, so their basketball experience involves the orange-ish, the α-size-ish, and the roundish characters. However, the planets are fundamentally different from Earth in the following ways:

> *Doubled Earth*: Dally is a normal perceiver on the planet where everything is doubled in size with respect to the things on Earth. On Doubled Earth, the standard meter stick in Paris is 2 meters long, the Taipei 101 is 1016 meters tall, and even Dally is twice as big as Sally. Doubled Earthlings perceive things in a way such that their meterstick experience involves the 1-meter-ish character, their Taipei 101 experience involves the 508-meter-ish character, and their experience of Dally involves the size character that is phenomenally identical to the size character involved in Earthlings' experience of Sally. Similarly, even though basketballs on Doubled Earth are twice as big as basketballs on Earth (call it β-size), Dally's basketball experience still involves the α-size-ish character instead of the β-size-ish character. (Case adapted from Thompson, 2010.)

> *El Greco Earth*: Gally is a normal perceiver on the planet where everything is "vertically stretched" to be twice as tall as things on Earth. The vertical stretching is dynamic on this planet. A person who is 183cm tall (from head to toe) when lying on the floor becomes 366cm tall when standing. Dice that have an equal chance for each face to face up are cuboid-shaped, balls that roll are elliptical, etc. But El Greco Earthlings perceive things in a way such that their experience of the person involves the 183cm-ish character (from head to toe), whether she is lying or standing. Their dice experience involves the cube-ish character, and their ball experience involves the roundish character. Although basketballs

on El Greco Earth are elliptical, Gally's basketball experience still involves the roundish character but not the elliptical-ish character. (Case adapted from Thompson, 2010; see also Hurley, 1998.)

Similar to CE1 and CE2, this section argues for two things:

SE1 *Doubled Earth* and *El Greco Earth* are possible, and

SE2 Neither Dally nor Gally suffers from misperception in size or shape.

SE1: Dally's and Gally's size and shape perceptions are not counterintuitive. We can imagine that if Earthlings went to Doubled Earth or El Greco Earth and wore goggles that halve or flatten things, respectively, they would not experience anything differently than they do on Earth. In this way, neither is it impossible that Dally perceives phenomenally the same size as Sally wearing size-altering goggles perceives on Doubled Earth, nor that Gally perceives phenomenally the same shape as Sally wearing shape-altering goggles perceives on El Greco Earth. Thus, it is not unreasonable to say that Sally and her phenomenal duplicates share the same phenomenal character while perceiving objects in different sizes and shapes. Perceptual variation in size and shape is indeed possible unless there is a strong reason against the possibility.

Some might reject the above reasoning on the grounds that spatial properties are not perceiver-dependent but perceiver-independent properties. The objection derives from the distinction of perceptible properties between primary and secondary qualities. According to Locke, the primary qualities of an object are "properties which the object possesses independent of us—such as occupying space, being either in motion or at rest, having solidity and texture". In contrast, the secondary qualities are "powers in bodies to produce ideas in us like color, taste, smell, and so on that are caused by the interaction of our particular perceptual apparatus with the primary qualities of the object" (Uzgalis, 2022: §2.2). Roughly in contemporary terminology, primary qualities are *perceiver-independent* perceptible properties, while secondary qualities are *perceiver-dependent* perceptible properties. In Locke's list, color is a secondary quality, but spatial properties, including size and shape (or, in his own words, extension and figure), are primary qualities (Locke, 1689/1975: Book. 2, Ch. 8, §9–10).

Based on the distinction, one could argue against the possibility of *Doubled Earth* and *El Greco Earth*. Since size and shape are primary qualities, what Dally and Gally perceive must differ from what Sally perceives. Suppose that none of them misperceives. Dally's basketball experience must involve a size character that is different from Sally's basketball experience, and the shape character in Gally's basketball experience must differ from Sally's. This results in inconsistency in the case of *Doubled Earth* and *El Greco Earth*, where Dally, Gally, and Sally are supposed to have experiences with the same phenomenology. So, the two planets seem impossible.

In addition, the possibility of *Color-inverted Earth* sheds no light on the possibility of *Doubled Earth* and *El Greco Earth*. Recall that a way to vindicate the possibility of *Color-inverted Earth* is that the commonality of color experiences can be explained by perceiver-dependent properties, such as sensational properties (Peacocke) or phenomenal properties (Shoemaker). The strategy succeeds because color is a secondary quality. But, it does not apply to *Doubled Earth* and *El Greco Earth* because size and shape are primary qualities. Since what Dally and Gally perceive is categorically different from what Sally perceives, the commonality of Sally's and Dally's size experience and that of Sally's and Gally's shape experience cannot be explained by their representing a common perceiver-independent size or shape property. Hence, although perceptual variation in color is possible, perceptual variation in size and shape is still impossible.

I reply: The opponents' objection relies on the assumption that spatial properties are primary qualities. But the assumption is not undebatable. If spatial properties are secondary qualities (if Locke mistakes them as primary qualities), the objection no longer holds.

Berkeley famously questions Locke's distinction between primary and secondary qualities. He argues that primary qualities cannot be abstracted from secondary qualities, since we cannot conceive of primary qualities that are separated from any secondary qualities. Ultimately, all primary qualities collapse into secondary qualities (Berkeley, 1713/1979, 1734/1982; see also Dicker, 2011: Ch.8). For example, we cannot conceive of an extended body without giving it "some colour or other [secondary] quality which is acknowledged to exist only in the mind" (Berkeley, 1734/1982: §10). If this is the case, then primary qualities cannot exist apart from secondary qualities. Since secondary qualities are perceiver-dependent, primary qualities are also perceiver-dependent. Thus, all perceptible properties are secondary qualities.

I am sympathetic to Berkeley's argument, but I do not extend the discussion because it goes beyond the scope of my argument. Even if Berkeley's argument fails, it motivates further

arguments that some perceptible properties traditionally considered primary qualities are secondary qualities. I argue that spatial properties are one of them.

Following Berkeley, Foster (1994) argues that the spatial properties we perceive the external reality as having are not identical to the spatial properties of the perceiver-independent physical world. Suppose there are two spaces, one of which is the perceiver-independent physical space (call it *P*), while the other is the space we perceive external reality as having (call it *E*). Thanks to current science, we can safely suppose that *P* has certain distributions of spatial properties over its points and regions and that any change in *P*'s spatial distributions across time is governed by the physical laws of *P*. Thus, *P*'s spatial distributions determine *P*'s geometrical structure. Also, suppose that *E* has certain distributions of spatial properties over its points and regions. There are changes in *E*'s spatial distributions across time, and the changes are governed by certain law-like experienced regularities, which, in particular, are regularities that "[prescribe] how earlier qualitative conditions [i.e., a certain spatial distribution of mental qualities at a given time] and processes give rise to later ones" (*Ibid.*: 517). For example, if there is a regularity in *E* that resembles the gravitational law in *P*, then we always see a ball "dropping" every time when it is released in the air. In this way, *E*'s spatial distributions determine *E*'s geometrical structure.

Now, are *P* and *E* identical? Physical realists, who take the perceiver-independent physical world as metaphysically fundamental, insist that physical entities are the elements of the fundamental space we perceive external reality as having (i.e., *E*). So, according to their view, *E* has to be numerally identical to the perceiver-independent physical space (i.e., *P*) or, at least, for some exhaustive division of *P*, each component of *P* relative to that division is numerically identical with some component of *E* (*Ibid.*: 522; see also Tse, 2022).[8] . In contrast, Berkeleyan idealists, who think the physical metaphysically depends on the mental, deny such an identity. Foster defends the latter view, arguing as follows:

F1 *P* possesses its geometrical structure essentially.

F2 *E* possesses its geometrical structure contingently.

[8] Foster suggests that giving realists the choice of this second and weaker alternative is to leave room for cases in which *E*'s and *P*'s intrinsic topologies do not coincide (*Ibid.*: 522). But this possibility does not affect my argument. For ease of discussion, we can safely suppose that *E*'s and *P*'s intrinsic topologies do coincide.

F3 For any two spaces, if they are numerically identical, then they must also have the same modal features.

F4 *P* and *E* do not have the same modal features.

F5 Therefore, *P* and *E* are not numerically identical. (Foster, 1994: §IV; as cited in Tse, 2022: 4)

F1 is true by default. As the distribution of spatial properties of *P* determines *P*'s geometrical structure, *P* must possess its geometrical structure essentially. As Foster (1994: 523) explicates, "the network of physical geometrical relations holding between the points and regions of *P* are [*sic*] essential to the identities of these points and regions: the latter could not be the particular points and regions they are without standing in these relations. […] If a certain point is physically located within a certain larger region, it could not have been physically located elsewhere, and that if two regions are physically contiguous, they could not have been physically separated." In short, it would no longer be *P* if *P* had a different spatial distribution.

But for F2, "whatever its intrinsic nature, *E* could have existed (with the identities of all its component points and regions unchanged) with a relevantly different [law-like experienced regularity], yielding a different physical geometry or no physical geometry at all" (*Ibid.*: 524). To clarify, imagine two physical spaces P_1 and P_2 (see **Figure 3-1** for pictorial illustration). Within P_1 and P_2, there are two three-dimensional regions, R_1 and R_2, such that the geometrical distribution of R_1 and R_2 in P_2 is reversed with respect to the geometrical distribution of R_1 and R_2 in P_1. In P_1, the geometrical structure is continuous between R_1 and its surrounding region U_1 and between R_2 and its surrounding region U_2. So, if a person in P_1 moves in a straight line from the center of R_1 to the center of R_2, her trajectory will be starting from R_1, passing through U_1 and U_2 chronologically, and arriving at R_2. In P_2, by contrast, R_1 is surrounded by U_2, and R_2 is surrounded by U_1. However, the geometrical structure in P_2 is not continuous between R_1 and U_2 and between R_2 and U_1, such that anyone arriving at some point on the boundary of R_1 while moving inwards or outwards instantaneously changes her location to the corresponding point on the boundary of R_2. Likewise, when one arrives at the boundary of R_2 while moving inwards or outwards, she instantaneously changes her location to the corresponding point on the boundary of R_1. Hence, if a person in P_2 moves in a straight line from the center of R_1 to the center of R_2, her trajectory will

be the same as the person in P_1: She starts from R_1, passes through U_1 and U_2 chronologically, and then arrives at R_2.

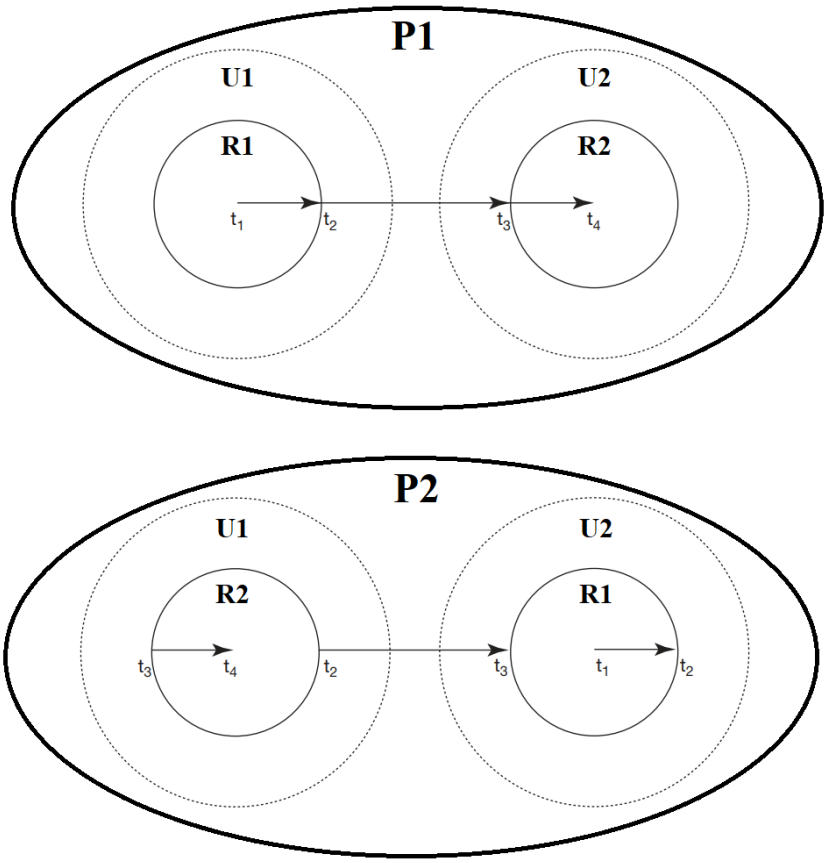

Figure 3-1. Two three-dimensional regions, R_1 and R_2, are reversed in two physical spaces P_1 and P_2. The arrows indicate the trajectory of the moving person in each space. The chronological order of trips is indicated by variable t_n. The figure is adapted from Foster (2008: 132).

P_1 and P_2 have distinct physical geometrical structures, as the former is continuous, but the latter is discontinuous. However, we cannot perceive such a difference. If a normal perceiver perceives the geometrical structure of P_1 as having E, she also perceives the geometrical structure of P_2 as having E. Suppose that one in P_1 normally perceives E, which has the nomological organization that R_1 is surrounded by U_1 and R_2 is surrounded by U_2. This is because our spatial perception is constrained by what we encounter through time. During a trip in P_1 from the center

of R_1 to the center of R_2, one typically perceives what she encounters chronologically (i.e., $R_1 \rightarrow U_1 \rightarrow U_2 \rightarrow R_2$) as spatially continuous. But the same perception happens in P_2. Due to the instantaneous change of points in its discontinuous geometrical structure, one's trip in P_2 from the center of R_1 to the center of R_2 appears chronologically the same as the trip in P_1. In this way, one in P_2 still normally perceives E. That is, one still perceives R_1 as surrounded by U_1 and R_2 as surrounded by U_2. Consequently, it is metaphysically equivalent for E to possess P_1 or P_2 as its physically geometrical structure. It is solely a matter of which physical space we contingently inhabit. Therefore, F2 is true (Foster, 1994: 517–8).

F1 and F2 imply F4. F3 is a criterion for identity. F3 and F4 imply F5: P and E are not numerally identical. The conclusion suggests that the spatial properties we perceive external reality as having are perceiver-dependent, secondary qualities. These spatial properties are not properties inherent in the physical world. Instead, they contingently relate to properties of the physical world.

Crucially, F5 justifies the possibility of *Doubled Earth* and *El Greco Earth*. First, it is not unreasonable for Dally and Gally to have size or shape experience that is phenomenally identical to Sally's experience. They can perceive the same experiential space E that contingently possesses different physical geometrical structures: Sally's E possesses Earth's physical geometrical structure (P_1); Dally's E possesses Doubled Earth's physical geometrical structure (P_2); and Gally's E possesses El Greco Earth's physically geometrical structure (P_3). Second, the physical geometrical structures of Doubled Earth and El Greco Earth are not internally inconsistent. Everything on Doubled Earth is physically scaled up twice as big as those on Earth, but Doubled Earth's physical geometrical structure is identical to the Earth's structure. On El Greco Earth, although only the vertical dimension of the physical space is scaled up, its topological structure is still the same as Earth's. These differences by no means hinder Dally and Gally from having E. Thus, *Doubled Earth* and *El Greco Earth* are indeed possible. SE1 is true.

My argument for SE2 resembles that for CE2. Stage 1: Sally's size and shape experiences are veridical. Stage 2: If Sally's size and shape experiences are veridical, then both Dally's size experience and Gally's shape experience are veridical. I briefly present the argument to avoid repetition, as the possible objections and responses remain the same.

Stage 1: Sally's size and shape experiences are veridical if they satisfy the lenient criterion of veridicality. First, they navigate successful interactions with the environment. Perceiving a basketball as having a specific size and shape enables Sally to detect the ball's motion more

efficiently. Second, they make relevant talks intelligible. When Sally describes a basketball by ascribing its size and shape that she perceives it as having, her friends will understand her because they also perceive the ball in the same way. Third, part of Sally's knowledge is justified by her size and shape experiences. At least, her knowledge that the basketball has a specific size and shape is partly derived from her size and shape experiences of the ball. Therefore, according to lenient veridicality, Sally's size and shape experiences are veridical.

Rejecting neither-ism: It requires strict veridicality, whose standard is too high to be included in lenient veridicality.

Stage 2: If Sally's size and shape experiences are veridical, so are Dally's and Gally's. Dally's size experience and Gally's shape experience are veridical if they satisfy the lenient criterion. Dally's size experience yields successful interactions with the environment, as she has no difficulty locating and catching the ball by perceiving it as α-sized. The experience also makes sense to other Doubled Earthlings because they all perceive it as α-sized. They understand each other when one says "The basketball is α-sized!" Moreover, when they claim certain knowledge, part of its justification may still be based on common perceptual features. Dally knows that the basketball is α-sized because that is what Doubled Earthlings perceive the ball as having. Consider Gally. Gally's shape experience facilitates her actions. Seeing the basketball as round may make Gally decide to wait for the ball to roll out of bounds. The experience makes sense to other El Greco Earthlings, for they have a common shape experience about the ball. The common experience also grounds El Greco Earthlings' common knowledge about the ball. All in all, since there is no difficulty for Dally's size experience or Gally's shape experience to satisfy the lenient criterion, they are veridical experiences. SE2 is true: Neither Dally nor Gally suffers from misperception in size or shape.

Rejecting one-ism: It must make a categorical difference in perceptual veridicality between Earthlings and Doubled Earthlings or between Earthlings and El Greco Earthlings. But there is probably no candidate to undergird such a difference.

To summarize: SE1 and SE2 are both true, so perceptual variation without misperception in size and shape is possible. With perceptual variation without misperception in color, I believe it suffices to generalize the phenomenon to the (visual) perception of other low-level properties. Therefore, PV1 is true, i.e., perceptual variation without misperception is possible. The following section turns to PV2.

3.4. Veridical Perception in Virtual Reality

The above arguments involve *inter*-subjective perceptual variation without misperception, i.e., perceptual variation across different perceivers. *Color-inverted Earth*, *Doubled Earth*, and *El Greco Earth* (the *Three Planets*, for short) all belong to this category. This section extends the inter-subjective variation without misperception to the *intra*-subjective variation without misperception, i.e., perceptual variation without misperception *within* the same perceivers. I argue that VR experiences fall under the latter category.

Today's VR technology allows users to experience (primarily through vision) various virtual things. In a virtual surgery, one may perceive a simulated organ as bleeding. In a virtual gallery, one may perceive colorful patches flying around the room. In an e-sports competition, one controlling an avatar may interact with teammates by perceiving other avatars' actions. When immersed in VR scenarios, the users may perceive the virtual as being real and be unable to distinguish the virtual from the non-virtual because the ways the virtual appears are so similar to the ways the non-virtual appears (Slater & Sanchez-Vives, 2016). Such a phenomenon leaves room for a subject to have the same phenomenal experiences in a virtual environment and a corresponding non-virtual environment.

Consider a case that combines the key features of the *Three Planets*:

> *Virtual Basketball.* Sally is a highly experienced VR user. Throughout her life, she spent a lot more time using VR than in her non-VR daily life. She is used to being in any virtual environment through advanced VR equipment sets that manipulate all her exteroception synchronously. The motion trackers track Sally's bodily movements and synchronously produce feedback to the virtual environment, so that her virtual avatar makes a corresponding movement at every moment when Sally moves her physical body. The virtual environment is so vivid that Sally cannot distinguish the virtual things from the corresponding non-virtual things. In other words, through introspection, Sally cannot distinguish between her VR experience and her corresponding non-VR experience. Now, a program generates a real-time counterpart of the non-virtual world, with which Sally is very familiar. In one scenario, the program generates a basketball on a court. Like other

users, Sally perceives the virtual basketball as orange, α-sized, and round, just like how she perceives an ordinary basketball in a non-virtual court.

Moreover, I argue that:

VE1 *Virtual Basketball* is possible in practice, and

VE2 *Virtual Basketball* does not involve misperception.

Here is my argument for VE1. *Step 1*: If the *Three Planets* are metaphysically possible, *Virtual Basketball* is metaphysically possible. *Step 2*: If *Virtual Basketball* is metaphysically possible, it is possible in practice.

Step 1: The difference between *intra*-subjective and *inter*-subjective variation is whether perceptual variation occurs *within* the same subjects or *across* different subjects. The *Three Planets* are cases of inter-subjective variation, where the same type of perceptual experience in different groups of subjects is reliably produced by distinct types of external properties in their respective environments. In contrast, *Virtual Basketball* is a case of intra-subjective variation, where a single subject has phenomenally identical experiences that are reliably produced by different external properties under different circumstances. A virtual basketball and an ordinary basketball have different perceptible properties, so the properties that reliably produce Sally's VR basketball experience (*VE*) are supposedly different in kind from the properties that reliably produce Sally's corresponding non-VR basketball experience (*N-VE*). Also, since Sally cannot distinguish between *VE* and *N-VE* via introspection, we can safely assume that *VE* and *N-VE* are phenomenally identical.[9]

Granted the possibility of the *Three Planets*, it is an easy move to the possibility of inter-subjective variation. Since different physical properties can reliably generate phenomenally identical experiences in different subjects, it is not unreasonable that these properties can reliably generate phenomenally identical experiences *within a single subject*. The intra-subjective variation

[9] This assumption is not universally accepted. For instance, a relationalist of perceptual experience does not agree. But in order to make *Virtual Basketball* a case parallel to the *Three Planets*, the thought experiment targets those accepting the common kind assumption of perceptual experience. Since they would naturally agree that subjective indistinguishability entails phenomenal identity, I set aside the controversy and assume that *VE* and *N-VE* have the same phenomenology.

in *Virtual Basketball* lies in the fact that *N-VE* comes from the group of experiences that are reliably produced by *virtual properties*, whereas *VE* comes from the group of experiences that are reliably produced by *physical properties*. Although a virtual basketball is not a physical object, the causal role it plays in Sally's perceptual machinery is the same as that of ordinary basketballs. In particular, it serves as a distal stimulus that causes the receptors on Sally's retina to activate and then generates *VE* through the visual process. Such a causal pathway is the same as the generating process of *N-VE*. With this respect, I find no hindrance for *Virtual Basketball* to be possible, even if it involves the perception of properties of heterogeneous kinds.

Step 2: Not only is *Virtual Basketball* metaphysically possible, but it is also possible *in practice*. Empirical cases of intra-subjective variation suggest that *Virtual Basketball* is actually possible. For instance, color metamerism (Hardin, 1988; Pautz, 2014) is a widespread case that different distributions of surface reflectance profiles reliably produce the same type of color experience in one subject.[10,11] Scientific evidence also shows that objects with different physical shapes reliably cause the same type of shape experience in one subject (Wagner, 1985; see also Green & Rabin, 2020). Moreover, given technological advancements, we can be confident that *Virtual Basketball* to be realized in the future. The high fidelity of VR nowadays has already made many users immersed in virtual scenarios so that they cannot easily distinguish between virtual scenarios and non-virtual environments. Once the technology further improves in the future, it is likely to realize the scenarios that appear all the same as non-virtual environments. By that time, *Virtual Basketball* will come true. Thus, *Virtual Basketball* is indeed possible in practice.

VE2: Provided that *N-VE* satisfies the lenient criterion of veridicality so that it is veridical, *VE* is also veridical if it satisfies the criterion.

First, *VE* facilitates successful interaction with the world. The color, size, and shape that Sally perceives the virtual basketball as having help Sally detect, locate, and track the ball in the virtual environment. If the ball and the visual background appear to have different colors, Sally will track the ball's movement more efficiently. If Sally sees the ball being passed to her, she may

[10] Specifically, color metamerism is based on the fact that human's three types of chromatically-sensitive retinal receptors (viz., long- (L), medium- (M), and short- (S) wavelength cone cells) jointly produce the response of receptors in the cone cells. Since the wavelength ranges to which these receptors are sensitive largely overlap, the same receptor response can be produced by light composed of distinct combinations of wavelengths. In effect, wavelengths reflecting different sets of surface reflectance properties reliably produce the same type of color experience.

[11] In fact, metamerism occurs not only in vision but also in audition (Helmholtz 1877/1954; Schouten 1940), olfaction (Margot, 2009), taste (Laffitte et al., 2017), and even haptic perception (Shockley et al., 2004), suggesting that color metamerism is not a special case in human perception.

get prepared to receive an appropriate tactile sensation by calculating the "speed" of the ball (i.e., by calculating the change rate of the ball's size in her visual field). Her perceiving the ball as being round prevents her from catching the ball in the same way as her catching an American football. Without those properties Sally perceives the virtual basketball as having, she would not interact with the ball so successfully.

Second, Sally's talk about the perceived color, size, and shape of the basketball is intelligible to other experienced VR users. Sally is supposed to be a user with a normal perceptual mechanism. In a particular virtual scenario, many phenomenal characters involved in *VE* and others' basketball experiences will be the same, including the orangish character, the α-size-ish character, and the roundish character. Given that commonality, Sally's ascription of these characters in her talk about basketball makes sense to other users. For example, others will understand her meaning when she utters "The orange ball is round!"

Third, the experiential basis of *VE* provides grounds for common knowledge. Given the existence of the virtual basketball, it is natural to ask about what properties the ball possesses. Such an inquiry also solidifies our knowledge of the ball. Sally's utterance "The orange ball is round!" may qualify as knowledge because of its reliable basis of *VE*. Specifically, with the perceptual condition being fixed, normal perceivers in Sally's perceptual community will share the orangish and roundish characters with Sally when they see the virtual basketball. If they perceive a virtual object that does not appear orange or round, they will know it is not a virtual basketball. They will probably agree with the common knowledge that the virtual basketball is orange and round, as they will also make a similar utterance while describing the ball's color and shape. Hence, the knowledge about virtual basketball is (at least partly) grounded in the commonality of the associated experiences.

Since *VE* satisfies the lenient criterion, it is leniently veridical. Together with the fact that *N-VE* is leniently veridical, VE2 is true. The criterion also makes sense of the falsidical VR experiences. Suppose Sally's VR equipment is malfunctioning, and she is not experiencing the colors or shapes of the virtual basketball as usual. Seeing the ball as being blue or cubical (say), she might probably say "Oh I misperceive the ball!" This suggests that a certain standard of veridical perception is met in her normal VR experiences (e.g., *VE*) and fails to be met this time. I believe the lenient criterion is just that standard. *VE* is leniently veridical because it satisfies the relevant conditions. But the bluish-cum-cubical experience does not satisfy them. It does not make

sense to relevant talks based on the altered experience because it is abnormal to the community. No other experienced users would understand what it means by "the blue and cubic basketball". Nor does the abnormal experience ground relevant knowledge, as the knowledge is underlain by what normal subjects in her community perceive the ball. So, according to the lenient criterion, the bluish-cum-cubical experience is falsidical.

Therefore, we have good reasons to believe that VE2 is true, i.e., *Virtual Basketball* does not involve misperception. With VE1, I conclude that *Virtual Basketball* is an instance of perceptual variation without misperception. Since there is no apparent obstacle to generalizing *Virtual Basketball* to other cases regarding (visual) perceptual experience in VR, many VR experiences are leniently veridical due to perceptual variation without misperception across VR and non-VR environments, i.e., PV2 is true. Consequently, given the justification of PV1 and that PV3 follows from the definition of virtual veridicalism, PV1 – PV3 together imply PV4: Virtual veridicalism is true.

3.5. Veridical Perception in Mixed Reality

Virtual veridicalism concerns perceptual experience in a virtual world. As I defined in Chapter 1, a virtual world is a shared, coherent simulated scenario. However, the extent to which a scenario is simulated comes in degrees. Presumably, we say a scenario is *perfectly* simulated when all entities that appear to the observers are virtual entities. We also say a scenario is *partially* simulated when some entities that appear to the observers are virtual entities and others are physical entities. Cases of perfectly simulated scenarios include VR games like *SUPERHOT VR* and *Beat Saber*. Cases of partially simulated scenarios are like *Pokémon GO* and *Angry Birds AR: Isle of Pigs*. We can also say that a virtual world is perfectly or partially virtual, depending on the degree of simulation. To reiterate my definition in Chapter 1, a virtual world is a *perfectly virtual world* when all entities that appear in the world are virtual entities, whereas it is a *mixed world* when some entities that appear in the world are virtual entities and others are physical entities.

The previous section defended virtual veridicalism by justifying *Virtual Basketball* as an instance of perceptual variation without misperception. But *Virtual Basketball* is a case in a perfect virtual world, so one might be concerned about whether my defense encompasses VR experiences in mixed worlds. To address this concern, this section generalizes the scope of virtual veridicalism

from perfect virtual worlds to mixed worlds, arguing that many VR experiences are leniently veridical not only in perfect virtual worlds but also in mixed worlds.

In VR technology, it is a consensus that the virtuality of a virtual world is a matter of degree. According to Milgram and Kishino's (1994) well-known framework of VR technology, virtual reality (VR), augmented reality (AR), and mixed reality (MR) all refer to a fragmented area with a vague boundary within a Reality-Virtuality Continuum (RV continuum). The RV continuum concerns the degree to which the virtuality blends with the non-virtual environment (see **Figure 3-2**). One side of the continuum is a perfectly non-virtual environment where no virtuality is involved, *viz.*, the ordinary environment. The other side is a perfect virtual world where no non-virtual thing is involved. An arbitrary point within the continuum refers to a mixed world, i.e., a scenario in which the virtual is mixed with the non-virtual to a certain degree. A mixed world located closer to the perfectly non-virtual side is an augmented reality (AR) scenario since the major components of this scenario are non-virtual, and virtuality is augmented on it. The same reason applies to the definition of an augmented virtuality (AV) scenario, where the major components of this scenario are virtual, and non-virtual things are augmented on it. Many factors affect the degree of reality-virtuality mixture in a certain scenario, but the mixture can still be described in one dimension. To simplify the terminology, I refer to a scenario with a certain degree of RV mixture by its virtuality percentage. For instance, the V-0% scenario is a perfectly non-virtual environment (a.k.a. an ordinary environment), and the V-100% scenario is a perfect virtual world. All scenarios from the V-1% scenario to the V-99% scenario are mixed worlds. A V-100% scenario becomes a V-99% scenario if less virtuality is involved, and vice versa. The difference between a perfect virtual world and a mixed world is only a matter of the degree to which virtuality is involved.

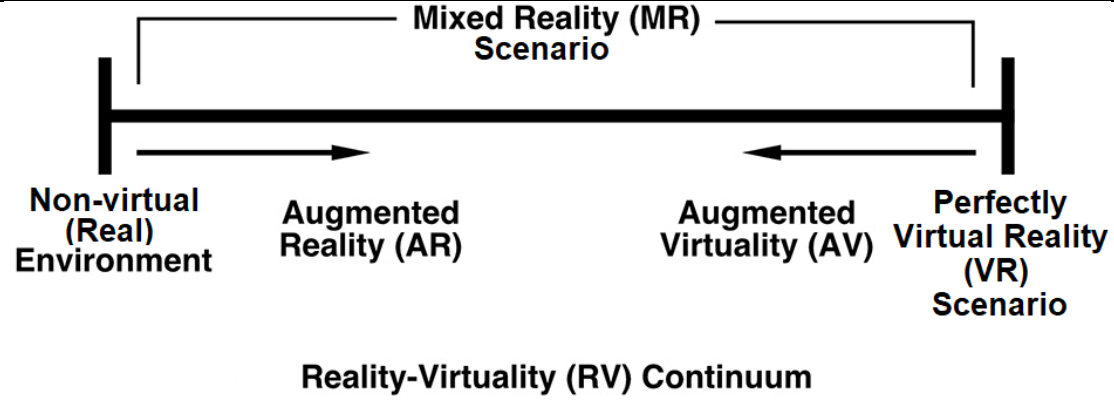

Figure 3-2. Milgram and Kishino's (1994) Reality-Virtuality Continuum. (Figure adapted from Skarbez et al., 2021: Fig. 1)

Based on the RV continuum, let me refine the story of *Virtual Basketball*. The advanced VR equipment sets that Sally wears are advanced glasses that cover the user's entire visual field. Ordinarily, the glasses function as common glasses. A user switches on the glasses when she wants to enter a virtual world. In doing so, the lenses block the array of lights from the outside, and the simulated scenario is shown on the lens. In *Virtual Basketball*, Sally perceives a V-100% scenario because the simulation completely screens off non-virtual entities through the lens. The glasses are so advanced that whenever Sally switches on the glasses, the appearance of the upcoming V-100% scenario is identical to that of the previous V-0% scenario. Hence, Sally cannot distinguish the V-100% scenario from the V-0% scenario just by their visual appearances.

Moreover, this kind of glasses can not only simulate a V-100% scenario but also make a partial simulation that blends virtual things into the non-virtual environment, thus creating a mixed world. Different people wearing the glasses can enter the same mixed world. When this function is switched on, the glasses block a significant proportion of the non-virtual environment and replace it with virtual counterparts, and the appearance of the new mixed world is identical to that of the previous V-0% scenario. We can make an adjustment to *Virtual Basketball* to fit this story:

> *Virtual Basketball in a Mixed World.* The advanced glasses block a proportion of the non-virtual environment with its virtual counterpart. The virtual part of the environment is so vivid that Sally cannot distinguish the virtual things from the corresponding non-virtual things. In the mixed world that the glasses simulate, Sally interacts well with virtual and non-virtual things without any problem. In one circumstance, Sally sees a basketball on a

court. This basketball is a virtual object. Like other users who wear these advanced glasses, Sally perceives the virtual basketball as orange, α-sized, and round, just like how she perceives an ordinary basketball in a non-virtual court.[12]

Given my argument in Section 3.4, it is natural to apply VE1 and VE2 to *Virtual Basketball in a Mixed World* and to justify the following claims:

ME1 *Virtual Basketball in a Mixed World* is possible in practice.

ME2 *Virtual Basketball in a Mixed World* does not involve misperception.

First, I have argued for VE1 (i.e., *Virtual Basketball* is possible in practice). If one accepts VE1, it is hard not to accept ME1. The main difference between *Virtual Basketball* and *Virtual Basketball in a Mixed World* rests on the difference in the imagined VR technology that simulates not only perfect virtual worlds but also mixed worlds. However, such a difference is not a gap in VR technology. Its high fidelity nowadays is thought to be capable of creating virtual worlds that are hardly distinguished from non-virtual environments. Given such techniques, it is not a problem to simulate mixed worlds. In fact, some products (e.g., the Apple Vision Pro or the Microsoft HoloLens) have already been capable of realizing a mixed world by adding virtual things into ordinary environments. As far as I can imagine, there is no principled obstacle for these products to realize a scenario described in *Virtual Basketball in a Mixed World*. So, ME1 is probably true.

According to the lenient criterion of veridicality, Sally's basketball experience in a mixed world should be deemed veridical. First, her basketball experience facilitates successful interaction with the mixed world. The color, size, and shape that Sally perceives the virtual basketball as having help Sally detect, locate, and track the ball in the mixed world. Without the properties Sally perceives the virtual basketball as having, she would not interact with the ball so successfully. Second, Sally's talk about the perceived color, size, and shape of the basketball is intelligible to other experienced VR users. Again, Sally is a normal VR user. Her talk about the virtual basketball

[12] Since this chapter concerns the experience of virtual things, I set the story in the way that, whatever the RV mixture ratio of the scenario is, the perceived basketball remains entirely virtual. This setting might be unable to defend virtual veridicalism in a very least augmented reality scenarios (e.g., V-0.001% ~ V-0.005% scenarios) due to the existence of the virtual basketball. We can ignore them because the range of scenarios to which my argument might be inapplicable is very limited.

makes sense to other users in the same mixed world. When she utters "The orange ball is round!" others will understand what she means. Third, Sally's basketball experience provides grounds for common knowledge. Through her experience, Sally knows that the basketball is orange, α-sized, and round, and such knowledge is common to those who perceive the basketball without error. Given that Sally's basketball experience satisfies the lenient criterion in the mixed world, the experience is leniently veridical. ME2 is true.

In sum, virtual veridicalism is true in both perfect virtual worlds and mixed worlds. Since virtual veridicalism concerns the perceptual experience of *virtual* objects and properties, and both worlds contain these entities, it is reasonable to generalize the scope of virtual veridicalism from perfect virtual worlds to mixed worlds.

3.6. Conclusion

To conclude, virtual veridicalism about low-level perceptible properties by the lenient criterion of veridicality is true. Having that perceptual variation without misperception is possible in color, size, and shape perception, I argued that experiences about color, size, and shape can be veridical in virtual worlds. Virtual veridicalism about low-level properties may help facilitate further arguments for high-level properties. Since the phenomenal character is essential to perceptual experiences, how it appropriately conveys the information from the environment is crucial to the prosperity of human beings. In this way, leniently veridical perception is indispensable to our daily lives. In addition, such indispensability also matters to our potential well-being in VR. If we expect to construct cohesive societies in VR in the future, whether we appropriately perceive the virtual worlds will soon be a pressing topic. For the present purpose, it suffices to conclude that we typically perceive virtual objects as having low-level properties that they really instantiate in virtual reality, including perfect virtual worlds and mixed worlds.

Admittedly, this chapter does not present the whole picture of virtual veridicalism. Despite its truth, the view needs to accommodate several theoretical questions. The first question concerns the relationship between VR experiences and virtual objects/properties. My argument in this chapter implies that certain virtual properties are represented by VR experiences so that VR experiences are veridical. Nevertheless, it does not identify the *nature* of these represented properties. In Chapter 2, I laid out an ontological picture of virtual worlds in which there are some candidates for qualifying as the represented properties. If the presented properties can be

precisified, virtual veridicalism will gain its metaphysical grounds. The second question would be about the contents of VR experiences. This chapter purposefully avoids discussing contents to establish a theory-neutral view from scratch. However, as different theories of content identify the nature of represented properties in different ways, and perceptual veridicality in turn depends on the nature of contents, a virtual veridicalist should clarify the constraints imposed by different theories of content. If this issue is clarified, we will have a firmer grasp of which theories cohere with virtual veridicalism and which do not. For virtual veridicalism to be a robust view, it must be able to answer these theoretical questions. The following chapters will investigate these questions in detail.

Chapter 4: Virtual Veridicalism and Representationalism

Chapter 3 concluded that virtual veridicalism is true: Many VR experiences are veridical. I defended the specific version that I call virtual veridicalism about low-level properties on lenient veridicality. That is, many VR experiences of low-level properties are leniently veridical. Also, these experiences are veridical in not only perfect virtual worlds (i.e., virtual worlds that *completely* screen off non-virtual things and then appear the same as non-VR environments) but also mixed worlds (i.e., virtual worlds that *partly* screen off non-virtual things and then appear the same as non-VR environments). But this is not the end of the story. Although many VR experiences are veridical, it is unclear *what they represent* in a virtual world. More specifically, what is the *content* of VR experiences? Also, due to perceptual variation across VR and non-VR environments, VR experiences and their corresponding non-VR experiences seem to appear the same to a subject. However, they also seem to represent different properties in different circumstances. Do they share the same content? Or do their contents differ? This chapter starts to look into the details of these issues.

As mentioned in Section 3.1, there are two ways to associate content with an experience: The *strict content* represents the property that appears in the experiences' phenomenology; while the *lenient content* represents the property that is partly determined by the appropriate contingent relationship between the experiences and the external world. In addition, I take the content of a perceptual experience to be its veridicality condition, i.e., the condition under which the experience is veridical. In this respect, there are also two senses with respect to which an experience is veridical: An experience is *strictly veridical* iff its strict content is true, while it is *leniently veridical* iff its lenient content is true. Since the specific version of virtual veridicalism defended in Chapter 3 rests on lenient veridicality, Chapters 4 and 5 look for a detailed picture of lenient contents that coheres with the version. I examine the coherence of this version within the representationalist framework. This chapter assesses currently prominent representationalist accounts, and Chapter 5 discusses my own account.

On the face of it, we can distinguish two types of properties of perceptual experiences: *phenomenal properties* and *representational properties*. For a perceptual experience to instantiate a phenomenal property is for it to exhibit a certain phenomenal character, which is a distinctive

what-it-is-likeness for the subject to undergo that experience. For a perceptual experience to instantiate a *representational property* is for it to represent a certain *content.*

Chapters 4 and 5 operate within a representationalist framework. *Representationalism* (or intentionalism) is roughly the view that a mental state's phenomenal properties are determined by its representational properties. Versions of representationalism can be distinguished along several dimensions. One dimension is the metaphysical strength of the relation between phenomenal and representational properties: *Identity representationalism* claims that phenomenal properties are *identical* to representational properties, while *supervenience representationalism* merely claims that the former *supervene* on the latter.[1] Representationalism also differs with respect to the relevant representational properties: *Pure representationalism* claims that phenomenal properties are identical to/supervene on *pure* representational properties, i.e., the properties of *representing a certain content*, whereas *impure representationalism* claims that phenomenal properties are identical to/supervene on *impure* representational properties, i.e., the properties of *representing a certain content in a certain manner*. For a specific content <red>, for instance, pure representationalism takes perceptually experiencing <red>, imagining <red>, and thinking about <red> as having the same phenomenal character; but impure representationalism allows that these mental states have different phenomenal characters since they involve phenomenal properties that differ in their manners of representation. Since these mental states presumably involve distinct phenomenal characters (if any), pure representationalism is hardly tenable, and most representationalists are impure representationalists. Impure representationalist views can be further divided into two types. *Intermodal representationalism* claims that phenomenal properties are identical to/supervene on impure representational properties of the form *representing C in manner M*, where *M* is the *same* manner of representation across all perceptual modalities. In contrast, *intramodal representationalism* claims that *M* is the manner of representation corresponding to each perceptual modality.[2] Since it is a commonly accepted view, I will adopt

[1] Many philosophers contend that the supervenience version is not a substantive position of representationalism. First, it is rarely denied, as every phenomenal experience is somewhat representational. Second, it is compatible with the existence of qualia, so even some anti-representationalists (e.g., Peacocke, 1983; Block, 1990, 1996) would agree with this version. Third, this version fails to *explain* the nature of phenomenal characters because supervenience is a superficial metaphysical relation that merely "*reports* a pattern of property covariation" (Kim, 1993: 167; italics mine).
[2] For details of distinctions of each version, see also Chalmers (2004), Bourget & Mendelovici (2014), and Bourget (2017).

the identity, impure, and intermodal version of representationalism, which I will call *representationalism* for short.

Representationalism is usually combined with a view of *phenomenal content* to explain the nature of phenomenal characters. Phenomenal content is usually understood as representational content that is determined by phenomenal character. *Russellian representationalism* takes phenomenal contents to be Russellian propositions, i.e., propositions made of specific objects and/or properties. In contrast, *Fregean representationalism* takes phenomenal contents to be Fregean propositions, i.e., propositions composed of modes of presentation (MoPs) of objects and properties. According to Fregeanism, MoP is an essential element of meaning of an expression. It is individuated by the cognitive significance of an expression and imposes the condition on extension (CoE), i.e., the condition that an entity in the world must satisfy in order to qualify as the expression's extension, depending on how the world turns out. In one version of Fregean representationalism, phenomenal content of an experience involves the MoP that imposes the CoE that picks out the property that normally causes experiences of that type (Chalmers, 2004, 2006; Thompson, 2009). Contents determine the conditions under which an experience is veridical or not. Different views of phenomenal content differ with respect to which properties phenomenal contents attribute, so they also differ with respect to the veridicality conditions of experiences.

My defense of virtual veridicalism focuses on low-level phenomenal characters such as colors, locations, and shapes. It is commonly assumed that perceptual experience has phenomenal contents of these characters. So, I assume that any experience with low-level phenomenal characters has phenomenal contents determined by these characters. It is straightforward to see that strict contents are phenomenal contents because they attribute the properties that appear in phenomenology. I also stipulate that lenient contents are phenomenal contents since, *prima facie*, experiences of the same sort always play the same representational role. I will discuss whether the lenient contents exactly count as phenomenal contents in Section 5.2.5, but it suffices to set this issue aside for now.

This chapter assesses three prominent representationalist accounts of phenomenal content to look for a theory of lenient content that coheres with the specific version of virtual veridicalism defended in Chapter 3. For terminological brevity, unless specified, I mean this specific version by "virtual veridicalism" in the rest of the dissertation. This chapter argues that none of the representationalist accounts being assessed cohere with virtual veridicalism, so they cannot be an

appropriate theory of lenient content for the latter view. Section 4.1 outlines the constraints embedded in virtual veridicalism. Section 4.2 rejects *physical Russellian representationalism*, according to which phenomenal content attributes physical properties. Section 4.3 rejects *dispositionalist Russellian representationalism*, the view that phenomenal content attributes dispositional properties. Section 4.4 rejects *Fregean representationalism*, which claims that phenomenal content involves modes of presentation of extensions.[3]

4.1. The Constraints of Virtual Veridicalism

My aim is to provide an account of the contents of perceptual experiences that is coherent with virtual veridicalism. In this section, I outline the constraints on such an account.

Since I am assuming representationalism, the account should satisfy its core thesis that an experience's representational contents determine its phenomenal characters. So, experiences with the same representational content have the same phenomenal character. In addition, as my defense in Chapter 3 focused on low-level phenomenal characters, I also assume that experiences have phenomenal contents (i.e., the representational contents that are determined by phenomenal characters) of these characters.

According to Chapter 3, many VR experiences are leniently veridical because they meet the lenient criterion, which is also met by ordinary, non-VR experiences. So, we have the first two constraints:

1. *Veridical VR Experience.* Many VR experiences are leniently veridical.
2. *Veridical Non-VR Experience.* Many non-VR experiences are leniently veridical.

Note that both VR experiences and non-VR experiences are perceptual experiences of human beings. Their difference lies in *where* or *in which environments* they are induced, resembling the difference between classroom experiences and garden experiences. In a classroom, our perceptual experiences are caused by things like blackboards, desks, etc. In contrast, our garden experiences

[3] According to Chalmers's (2004) taxonomy, representationalist accounts can be divided roughly into five categories. Four are Russellian accounts, namely physical Russellian representationalism, dispositionalist Russellian representationalism, projectivist Russellian representationalism, and primitivist Russellian representationalism. The other is Fregean representationalism. Since the primary concern of projectivist and primitivist Russellian representationalism is the strict contents, not the lenient contents, and virtual veridicalism (my version) concerns lenient contents, projectivist and primitivist accounts are irrelevant to my aim here.

are caused by things like flowers, trees, etc. Similarly, in a virtual world, VR experiences are caused by virtual objects and properties, while in an ordinary environment, non-VR experiences are caused by non-virtual objects and properties. It would be unreasonable to suppose a fundamental difference between VR and non-VR experiences.

Also, my defense in Chapter 3 is based on perceptual variation without misperception across VR and ordinary environments. The core is that two groups of token perceptual experiences of the same type share some phenomenal characters while being reliably caused by categorically different properties. In perceptual variation in VR users, a group of token experiences of a certain type are VR experiences, i.e., they occur in a virtual world, while the other group of token experiences of the same type are non-VR experiences, i.e., they occur in an ordinary environment. There are VR experiences of many types, and so are non-VR experiences. Following my terminology from above, I hereby define, more specifically, VR experience and non-VR experience as the *group* of token experiences of a certain type that occur in a virtual world and an ordinary environment, respectively. Also call a pair of VR experience and non-VR experience of the same experience type the *corresponding* experiences. Reddish VR experience is the corresponding experience of reddish non-VR experience, and vice versa. The VR experience is reliably caused by virtual properties, but the corresponding non-VR experience is reliably caused by non-virtual properties. So, we have two more constraints:

3. *Same Phenomenology*. A VR experience and its corresponding non-VR experience sometimes share some phenomenal characters.
4. *Different Reliable Cause.* The properties that reliably cause a VR experience are different in kind from those that reliably cause its corresponding non-VR experience.

At last, Section 3.5 generalizes the scope of virtual veridicalism to mixed worlds, suggesting an additional constraint that the above four constraints are satisfied in not only perfect virtual worlds but also mixed worlds. If an account satisfies the above four constraints only in perfect virtual worlds, it does not satisfy the fifth constraint:

5. *Scope. Veridical VR Experience*, *Veridical Non-VR Experience*, *Same Phenomenology*, and *Different Reliable Cause* are satisfied in both perfect virtual worlds and mixed worlds.

These five constraints exhaust what virtual veridicalism (i.e., the specific version defended in Chapter 3) requires. An account satisfying only the constraints 1– 4 coheres with virtual veridicalism only in perfect virtual worlds, and an account satisfying all five constraints coheres with virtual veridicalism in not only perfect virtual worlds but also mixed worlds. In this respect, this chapter assesses the coherence of virtual veridicalism with three representationalist accounts by examining whether they are resourceful enough to accommodate all the constraints.

4.2. Physical Russellian Representationalism

This section assesses physical Russellian representationalism, according to which the phenomenal properties of perceptual experiences are determined by Russellian representational contents involving the attribution of physical properties to represented objects (Dretske, 1995; Lycan, 1996; Tye, 1995, 2000). I argue that this account cannot accommodate the virtual veridicalist constraints, so it fails to cohere with virtual veridicalism.

Physical Russellian representationalism is committed to three claims:

Representationalism. An experience's representational contents determine its phenomenal characters.

Russellianism. The representational content of an experience is composed of specific objects and properties.

Tracking Thesis. An experience's representational content is fixed by what it tracks, i.e., by the objects and properties that the experience detects, carries information about, or otherwise correlates with.

Physical Russellian representationalism's commitment to *Representationalism* and *Russellianism* is indisputable. Different versions of this account characterize the tracking relation in different ways. Dretske (1995) considers the tracking relation as an information indication relation. In his view, an experience *E* represents the external property *F* iff *E* has the function of indicating (i.e.,

providing information about) *F*, where the function is derived from its evolutionary history. Lycan (1996) also endorses a similar evolutionary approach to the tracking relation. Tye (2000) considers the tracking relation as causal covariance under "optimal conditions". Accordingly, an experience *E* of a creature *c* represents that *P* just in case "if optimal conditions were to obtain, [*E*] would be tokened in *c* iff *P* were the case; moreover, in these circumstances, [*E*] would be tokened in *c* because *P* is the case" (Tye, 2000: 136). All in all, they reduce the representational relation to a sort of tracking relation, and the representational content of an experience is fixed by what it tracks (see also Bourget & Mendelovici, 2014: 214).

The three doctrines together entail the following thesis:

> *Tracking Phenomenal Externalism*. The phenomenal content of an experience attributes the external property tracked by that experience.

Representationalism, *Russellianism*, and the *Tracking Thesis* together entail the claim that the tracked external property itself determines what it is like for us to experience it. Hence, any physical Russellian representationalist must endorse *Tracking Phenomenal Externalism*.

Still, *Tracking Phenomenal Externalism* does not suffice for physical Russellian representationalism. The thesis entails that the phenomenal content of an experience *attributes* the external property tracked by the experience. Nevertheless, the external properties may or may not be physical properties. If the properties tracked by an experience are physical properties, the proponents would predict that its phenomenal contents attribute physical properties. Although what an experience tracks need not be physical properties, most philosophers accepting *Tracking Phenomenal Externalism* implicitly identify the former with the latter (Dretske, 1995; Lycan, 1996; Tye, 1995, 2000), since they aim to integrate their view into the scientific image of the world by accounting for perceptual representation in terms of physical notions. So, they accept an extra claim:

> *Identification*. The properties tracked by an experience are specific physical properties.

The four doctrines together constitute physical Russellian representationalism:

> *Physical Russellian Representationalism*. The phenomenal content of an experience attributes the external *physical* property tracked by that experience.

In other words, an experience represents a particular physical property if the experience tracks it. The representation relation also determines an experience's veridicality conditions: An experience is veridical iff its object instantiates the physical property it represents as having. Suppose that reddish experiences track the physical property of *having a surface reflectance profile that reflects light with a wavelength of 700nm* (C_{700}, for short). Physical Russellian representationalism predicts that the former represents the latter: A reddish experience is veridical iff its object instantiates C_{700}.

It is widely held that *Physical Russellian Representationalism* conflicts with the possibility of spectrum inversion without illusion.[4] Spectrum inversion is a possibility in which two groups of perceivers have systematically inverted color experiences when looking at the same objects (Locke, 1689/1975: II, xxxii, 15; see also Byrne, 2020: §3.2). Suppose Jack and Jill are normal perceivers of each group. Jack perceives grass as phenomenally green and ripe tomatoes as phenomenally red. In contrast, Jill perceives grass as phenomenally red and ripe tomatoes as phenomenally green. Presumably, it is hard to say that either Jack or Jill suffers from color illusion, as "there is no sense in which the experience of one of them is more or less true to the objective nature of what is experienced, namely tomato [or grass], than the experience of the other" (Shoemaker, 1994: 27). But if both Jack's and Jill's color experiences are granted veridical, then such a case is incompatible with *Physical Russellian Representationalism*. By assumption, the reddish experience Jack has when viewing a ripe tomato is phenomenally identical to the reddish experience Jill has when looking at grass. If Jack's reddish experience is veridical, it is because he represents the tomato as having the property that the tomato instantiates. The property, in this case, is C_{700}. If so, Jill's reddish experience must also represent the same property: She must represent the grass as having C_{700}. However, since the grass never instantiates C_{700}, Jill's reddish experience can never be veridical. Similarly, if Jill's reddish experience is veridical seeing grass, it must represent the grass as having the property that the grass instantiates, i.e., the physical property of *having a surface reflectance profile that reflects light with a wavelength of 550nm* (C_{550}, for short). If so, Jack's reddish experience can never be veridical because ripe tomatoes never instantiate C_{550}. In short, *Physical Russellian Representationalism* predicts that either Jack or Jill misperceives.

[4] In fact, Nida-Rümelin (1996) argues that there are probably actual cases of spectrum inversion.

For similar reasons, *Physical Russellian Representationalism* is incompatible with virtual veridicalism. Consider reddish experiences in the same subject that occur in different possible situations. A group of her token reddish experiences occurs in a situation where she is in an ordinary, physical environment, and this group of experiences is reliably caused by C_{700}. Another group of her token reddish experiences occurs in a situation where she is regularly immersed in a virtual world, and that group of experiences is reliably caused by a virtual property (call it VC_{700}). Call a token experience from the first group E_P and a token experience from the second group E_V. As E_P and E_V come from the groups of reddish experiences that are reliably caused by different properties, the E_P/E_V pair satisfies *Different Reliable Cause*. Also, both E_P and E_V are reddish experience, so they have the same phenomenal character. *Same Phenomenology* is satisfied. Depending on the relevant tracking relation, reddish experience might track C_{700} or VC_{700}. Since VC_{700} is not a physical property, if reddish experience tracks VC_{700}, then this contradicts *Physical Russellian Representationalism* that experiences represent physical properties. If, instead, reddish experience tracks C_{700}, then E_V misrepresents in virtual worlds, and this contradicts virtual veridicalism because *Veridical VR Experience* is never satisfied.

A physical Russellian representationalist might raise two immediate responses. First, she might say that E_P belongs to one experience type that tracks C_{700} while E_V belongs to another experience type that tracks VC_{700}, and also insist that VC_{700} is a distinctive physical property. If so, E_P and E_V both can be veridical while preserving *Physical Russellian Representationalism*. However, even if VC_{700} is physical, this response is still incompatible with virtual veridicalism. In this case, E_P and E_V no longer share the same phenomenal character. *Same Phenomenology* is thus violated.

Second, the physical Russellian representationalist might respond the other way around. She might identify VC_{700} as the same physical property as C_{700}, since they reliably cause the same experience type. So, this is an independent reason for not taking E_P and E_V to track different properties. Instead, they both track and hence represent the same physical property. However, this response rejects *Different Reliable Cause*. Although it might harmonize physical Russellian representationalism with some versions of virtual veridicalism, it is inconsistent with the version I defended in Chapter 3. Thus, this response does not rescue the account for the present purpose.

If physical Russellian representationalists intend to cohere with virtual veridicalism while sticking to the presupposed four claims, a solution might be to argue that reddish experience tracks

and represents a certain Lockean secondary quality, namely, phenomenal redness, and to insist that such the secondary quality is "broadly physical". For instance, Cutter (2018) contends that phenomenal redness is broadly physical insofar as it can be fully explained by its *causal dispositions* together with its *intrinsic qualitative nature*. Specifically, current physical science can fully characterize the causal relation of phenomenal redness, such that phenomenal redness disposes its physical bearers, namely C_{700} and VC_{700}, to cause reddish experience. In addition, although current physics leaves unspecified the intrinsic qualitative nature of phenomenal redness, Cutter argues that it is merely a verbal dispute whether to call phenomenal redness "physical" or not because it still "serve[s] as the basis for the basic causal dispositions characterized by physics" (*Ibid.*: 47). So, there is no problem also to count phenomenal redness as a broadly physical property. If Cutter's solution is viable, physical Russellian representationalists will have room to say that E_P and E_V track and represent phenomenal redness both in ordinary environments and in virtual worlds.

I am sympathetic to the move to the secondary qualities.[5] But awkward consequences arise if they count as physical properties. Although physical science offers a structural and dynamic picture of the world, it cannot fully explain *what it is* for a physical property to give rise to a phenomenal type. A physical explanation of phenomenal redness must account for the commonality of E_P and E_V without appeal to phenomenal notions. The only explanation would be that redness is grounded in the disjunction of C_{700} and VC_{700}. This seems *ad hoc* because it does not account for *why* these physical properties, not anything else, commonly cause the reddish experience. Furthermore, even if the disjunctive explanation is reasonable, it is still questionable why a virtual property (i.e., VC_{700}) could ground a physical property (i.e., phenomenal redness), as virtual property is categorically different from physical property. Unless an *ad hoc* explanation is given, the solution from the secondary qualities does not save physical Russellian representationalism.

In sum, physical Russellian representationalism does not satisfy the constraints of virtual veridicalism, and I conclude that it is not consistent with virtual veridicalism, as defended in Chapter 3.

4.3. Dispositionalist Russellian Representationalism

[5] To forecast, I will propose an account involving the representation of higher-order properties in Chapter 5.

This section assesses dispositionalist Russellian representationalism, according to which the phenomenal property of perceptual experience is determined by Russellian representational property involving the attribution of dispositional properties. I focus on Shoemaker's (1994, 2000, 2001b, 2003, 2006) account, which is the most developed version in the literature. Unless specified, I mean Shoemaker's account by "dispositionalist Russellian representationalism". I argue that his account does not satisfy the constraints of virtual veridicalism, so it is inconsistent with the view.

Shoemaker's aim is to develop a version of Russellian representationalism that avoids the problems encountered by physical Russellian representationalism. Specifically, he aims to accommodate Russellian representationalism and the possibility of spectrum inversion without illusion. So, he commits to *Representationalism* and *Russellianism*:

> *Representationalism.* An experience's representational contents determine its phenomenal characters.
> *Russellianism*. The representational content of an experience is composed of specific objects and properties.

Also, he endorses the possibility of spectrum inversion without illusion. As discussed above, physical Russellian representationalism is incompatible with such a possibility. Shoemaker proposes a Russellian representationalist account to resolve this issue.

Shoemaker argues that color experience does not represent physical properties but *appearance properties*. At first pass, appearance properties in general, involve dispositions to cause experiences of a certain sort. According to Shoemaker, identifying the represented property in this way reconciles his above-mentioned commitments. Suppose Jack is looking at a ripe tomato and Jill is looking at grass. Since they both have reddish experiences, they share the phenomenal content that attributes the same property to the ripe tomato and the grass, respectively. For both of their reddish experiences to be veridical, the ripe tomato and the grass must have the represented property. The ripe tomato and the grass do not share physical properties of having the same surface reflectance profile, so the represented property cannot be these physical properties. Rather, Jack's and Jill's reddish experiences represent the ripe tomato and the grass, respectively, as having the same appearance property, say, the property of *being disposed to cause reddish experience*. Assuming that the ripe tomato and the grass do have such a property (otherwise, Jack and Jill

cannot perceive them as red), Jack's and Jill's reddish experiences are both veridical. Furthermore, this account violates neither *Representationalism* nor *Russellianism* nor the possibility of spectrum inversion without illusion. It seems tenable at first glance.

Similarly, Shoemaker's account seems to accommodate the constraints of virtual veridicalism. Recall the E_P/E_V pair, where E_P is a token of a group of reddish experiences that is reliably caused by C_{700} in an ordinary environment, while E_V is a token of another group of reddish experiences that is reliably caused by VC_{700} in a virtual world. First, E_P and E_V are reddish experiences. They have identical phenomenal character. *Same Phenomenology* is satisfied. Second, E_P and E_V represent the same appearance property (e.g., *being disposed to cause reddish experience*), so it is compatible with the fact that they are reliably caused by distinct physical properties. Shoemaker is explicit about this point, saying that "what must be true of different tokens of the same [experiential] type is that they share a certain causal role […] and so for the possibility that tokens of an experiential type might be physically heterogeneous" (Shoemaker, 2000: 267). Distinct physical properties can play a certain causal role in producing the same type of experience. Thus, his account satisfies *Different Reliable Cause*. Third, since E_P and E_V represent the same property in ordinary and virtual environments, respectively, and the property can be instantiated in both environments, E_P and E_V can be veridical. That is, *Veridical VR Experience* and *Veridical Non-VR Experience* are satisfied. Therefore, dispositionalist Russellian representationalism seems to cohere with virtual veridicalism.

However, although this account seems to accommodate the constraints of virtual veridicalism, it makes misrepresentation by experiences impossible. Since every object has the disposition to cause an experience of any sort, if any experience represents its object as having such a dispositional property, then it must be veridical. Suppose that a virtual ball instantiates VC_{700}, which is supposed to generate reddish experience. Even if one perceives the ball as phenomenally green or blue or other colors, Shoemaker's account still predicts that her experience is veridical because the ball is disposed to cause that experience (otherwise, she cannot perceive it

in this way). As a result, this account is unacceptable due to the impossibility of misrepresentation.[6,7]

Hence, the nature of the appearance properties should be more qualified to avoid the impossibility of misrepresentation. Shoemaker (2001b: 254 ff.) considers three kinds of appearance properties as the potential dispositional properties represented in experiences. They are *occurrent appearance properties* (i.e., the properties of *causing an experience of a certain sort in a particular perceiver*), *dispositional appearance properties* (i.e., the properties of *being disposed to cause an experience of a certain sort in some kinds of observers under certain circumstances*), and *higher-order dispositional appearance properties* (i.e., the properties of *having one of the dispositional appearance properties that can manifest itself in the instantiation of a particular occurrent appearance property*).

However, Thompson (2007) nicely argues that none of these appearance properties qualify as a correct theory of phenomenal content. First, the occurrent appearance properties simply do not escape from the impossibility of misrepresentation. Assuming that an experience is caused at all, something will always have the occurrent appearance property the experience represents it as having whenever it gets represented (*Ibid.*: §4). There is no room for an experience to represent its object as having the occurrent appearance property it does not have.

Second, the dispositional appearance properties. An inherent problem threatens the internal consistency of this route. This problem concerns the individuation of the dispositional appearance properties: Since "subject groups" and "perceptual circumstances" are involved in the qualification of dispositional appearance properties, the individuation of the dispositional appearance properties has to be relativized with respect to them.

[6] Some might wonder whether one could be in a situation where we experience an object as red, but the object plays no role in making us see red, such that redness is just a property of our experience but not of the perceived object. In other words, we "project" the property that only appears in experiences onto the object. If so, our experience does misrepresent the object as having properties that it does not have. In response, this is indeed a possibility. But such a possibility is ruled out from Shoemaker's project because in his view the properties we experience the object as having "really do belong to the external things in which we perceive them as being instantiated" (2000: 252). For the similar reason discussed in my footnote 4, I set aside the projectivist alternative to the dispositionalist view.

[7] Some might wonder whether the problem of impossible misrepresentation here can be resolved by adding something like "normal conditions" into the characterization of the appearance properties, such that reddish experience represents something like the property of *being disposed to cause reddish experience under normal conditions*. If so, one misrepresents when the experience is abnormal. Indeed, I believe this solution will resolve the problem. In Chapter 5, I will adopt this route to propose my account. But since Shoemaker's account does not have resources to include "normal conditions" in phenomenal content, I set aside this possible solution in this section.

As Thompson (*Ibid.*: §5) argues, a dilemma arises regarding the relativization of "subject groups". In principle, the individuation of dispositional appearance properties can be relativized with respect to individual subjects or certain groups of subjects. If it is relativized to individual subjects, then no different subjects share the same phenomenal character. Specifically, if a dispositional appearance property is individuated with respect to Jack and Jill, respectively, then what Jack's reddish experience represents is the property of *being disposed to cause reddish experience in Jack*, whereas Jill's reddish experience represents the property of *being disposed to cause reddish experience in Jill*.[8] Since Jack's and Jill's reddish experiences represent different properties, their experiences are not phenomenally identical, contradicting the setting of spectrum inversion without illusion. Hence, dispositional appearance properties cannot be relativized to individuals.

Yet, if dispositional appearance properties are individuated with respect to certain groups of subjects, more problems arise. On one hand, if Jack and Jill are categorized into different groups, the same problem occurs. Jack's reddish experience represents the property of *being disposed to cause reddish experience in Jack's group*. In contrast, Jill's reddish experience represents the property of *being disposed to cause reddish experience in Jill's group*. Again, this solution fails because it implies a difference in Jack's and Jill's reddish phenomenology, and the same contradiction follows.

On the other hand, if Jack and Jill are categorized into the same group, some awkward consequences follow. Suppose their reddish experiences represent the property of *being disposed to cause reddish experience in Jack OR Jill.* Although Jack's and Jill's reddish experiences are both veridical, they are veridical even if they switch their perceived objects. In particular, it is predicted that Jack's reddish experience of grass and Jill's reddish experience of a ripe tomato are still veridical. As a result, this account wrongly classifies some apparent cases of misrepresentation as veridical.

By contrast, if Jack's and Jill's reddish experiences represent the property of *being disposed to cause reddish experience in Jack AND Jill*, their experiences can never be veridical. As Jack and Jill are color-inverted, no particular property can be disposed to cause the same experiential type in Jack and Jill at once. Such a conjunctive property is never instantiated, so

[8] I assume that Jack and Jill are in the same perceptual circumstance for convenience, so I get rid of the "under certain circumstances" part when specifying the represented property here.

Jack's and Jill's reddish experience cannot veridically represent. Categorizing Jack and Jill into the same group is untenable. Consequently, the dilemma imposes a serious problem upon the appeal to the representation of dispositional appearance properties. There is no good solution to relativizing "subject groups".

Similarly, since dispositional appearance properties also involve the qualification of "perceptual circumstances", their individuation has to be relativized with respect to this latter notion. In light of the problem of relativizing "subject groups", I argue that a similar dilemma arises when relativizing "perceptual circumstances", and such a dilemma makes dispositionalist Russellian representationalism inconsistent with virtual veridicalism.

Consider the E_P/E_V pair again. First horn: If the dispositional appearance properties are relativized to particular circumstances, then no pair of E_P and E_V represents the same property. Since ordinary and virtual environments are distinct perceptual circumstances, E_P represents the property of *being disposed to cause reddish experience in the ordinary environment*, but E_V represents the property of *being disposed to cause reddish experience in the virtual world*. Since phenomenal contents attribute these properties, the difference in their represented properties entails that E_P and E_V have distinct phenomenal contents and, in turn, different phenomenal characters. This horn violates *Same Phenomenology*, ruling out the possibility that E_P and E_V can be phenomenally identical.

Second horn: If the dispositional appearance properties are relativized to certain types of perceptual circumstances, the demarcation of the types matters. On the one hand, if ordinary and virtual environments are categorized into different types of circumstances, *Same Phenomenology* is not satisfied, either. This is because E_P represents the property of *being disposed to cause reddish experience in circumstances like the ordinary environment*, but E_V represents the property of *being disposed to cause reddish experience in circumstances like the virtual world*. Due to the difference in the represented property, E_P and E_V have different phenomenology. On the other hand, if ordinary and virtual environments are categorized into the same type of circumstances, then awkward consequences follow. If E_P and E_V represent the property of *being disposed to cause reddish experience in the ordinary AND virtual environments*, then we must reject both *Veridical VR Experience* and *Veridical non-VR Experience*. Because the nature of instantiated properties categorically differs in the ordinary environment (i.e., physical properties) and in the virtual world

(i.e., virtual properties), no objects can instantiate such a conjunctive dispositional property. So, neither E_P nor E_V can be veridical.

If E_P and E_V represent the property of *being disposed to cause reddish experience in the ordinary OR virtual environments*, *Veridical VR Experience* and *Veridical non-VR Experience* are both satisfied. However, this solution is pretty *ad hoc*, as it is hard to explain why there is such a type of property. First, VR could never have been invented. Were it the case, it would be strange to posit a property that is instantiated in a non-existent perceptual circumstance. Second, E_P and E_V cannot misrepresent. Since anything that has the potential to produce reddish experience (including E_P and E_V) has the dispositional appearance property, it is impossible for a reddish experience to misrepresent objects as having the property they do not have. Third, it is counterintuitive to explain the naïve VR users' perception. Presumably, when one uses VR for the first time and perceives something in the virtual world, her VR experience would probably go wrong because she has never encountered virtual objects and properties and knows nothing about them. But the solution implies that even though the subject has no idea about VR, the first encounter with virtual things produces veridical VR experiences, since the experience represents the same property as her non-VR experiences do. For these reasons, the appeal to dispositional appearance property does not offer a viable solution to make dispositionalist Russellian representationalism consistent with virtual veridicalism.

Before ending this sub-section, let's see whether the third kind of appearance properties, i.e., the higher-order dispositional appearance properties, provides a way out of the difficulty discussed above. Like the occurrent appearance properties, taking higher-order dispositional appearance properties as the properties being represented also makes misrepresentation impossible. Thompson (2007: §6) has argued for this point in detail. I will not rehearse his full argument here. But let me illustrate it in an extreme case. Suppose Jack's visual system is temporarily malfunctioning, so he perceives grass as phenomenally red. This is a case of misrepresentation. However, if phenomenal content attributes higher-order dispositional appearance properties, then Jack's reddish experience of grass is deemed veridical, since grass does have a higher-order property of having the dispositional appearance property that manifests itself when it causes reddish experience in Jack under the perceptual circumstance he is in. Therefore, it is not the case that perceptual experience represents higher-order dispositional appearance properties, either.

4.3.1. The Indexical Strategy to the Rescue?

Dispositionalist Russellian representationalism is problematic because of its internal tension regarding the relativization of "subject groups" and "perceptual circumstances" involved in the individuation of the dispositional appearance properties. This sub-section looks for possible solutions to this problem. I argue that the solution renders dispositionalist Russellian representationalism no more coherent with virtual veridicalism.

To resolve the problem of "subject groups", dispositionalist Russellian representationalists may adopt an indexical strategy to avoid the demand for relativization. For instance, Egan (2006a, 2006b) argues that the phenomenal content of perceptual experience is *self-locating content.* Typically, we take mental contents as *possible-worlds* propositions and describe such possible-worlds contents in terms of the *properties* that things are represented as having. By representing possible-worlds contents, one represents what the world is like. By representing John Perry as having the property of *being an American*, we pick out a set of worlds — all and only the worlds in which John Perry is an American. By representing David Lewis as having the property of *being an Australian* we pick out another set of worlds — all and only the worlds in which David Lewis is an Australian.

Self-locating contents, by contrast, are *centered-worlds* propositions. A centered world is a set-theoretic object that models a location within a world. Roughly, a centered world can be understood as a <world, <individual, time>> pair (Egan, 2006b: 107; see also Perry, 1979; Lewis, 1979). By representing self-locating content, one represents not only what the world is like but also one's location within the world.[9] I can represent a friend as *being nearby* or others as *being far away*; I can also represent an umbrella as *being to my left* and a bee as *being on my nose*. Specifically, when I represent a friend as *being nearby*, my self-locating content is the set of centered worlds in which a friend is near the center, i.e., it is all and only <*w*, <*i*, *t*>> pairs such that a friend is near *i* at *t* in *w*. Also, when you represent a friend as *being nearby*, your self-locating content is the same set of centered worlds. So, my representation <a friend is nearby> and your representation <a friend is nearby> share the same self-locating content.

[9] It is worth noting that "location" in this sense need not be geographical. Many self-locating contents such as <my hat is too tight>, <Fozzie is staring at me>, and <limes are disposed to cause greenish experience in me under illuminant circumstance> have nothing to do with geographical self-location (Egan, 2006a: 511).

If dispositionalist Russellian representationalism adopts the indexical strategy, the problem of relativizing "subject groups" can be dissolved. Since this strategy relativizes "subject groups" with respect to individuals indexically, the dilemma addressed by Thompson does not arise. As Egan suggests, the phenomenal content of an experience is the self-locating content that attributes the property (*F*) of *being disposed to cause an experience of a certain sort (E) in me under certain circumstances (C)*. The "*in me*" part helps pick out the set that includes all and only centered worlds in which the individual of the center is disposed to have *E* from objects instantiating *F* in *C*. So, when Jack sees a ripe tomato, he perceives it as having the property of *being disposed to cause reddish experience in me under a certain circumstance*. When Jill sees grass, she also perceives it as having the property of *being disposed to cause reddish experience in me under a certain circumstance*. Be they in the same perceptual circumstance, their reddish experiences share the same phenomenal content and represent the same property.

I believe the indexical strategy makes dispositionalist Russellian representationalism a promising account.[10] But even this strategy does not help the account resolve its inconsistency with virtual veridicalism. In particular, even assuming that the problem of relativizing "subject groups" is resolved, the problem of relativizing "perceptual circumstances" remains. As I argued above, it is the latter that makes dispositionalist Russellian representationalism inconsistent with virtual veridicalism. Given the indexical treatment of "subject groups", it is unreasonable also to relativize "perceptual circumstances" in terms of the same strategy because it will result in the impossibility of misrepresentation again. For instance, if reddish experience represents the property of *being disposed to cause reddish experience in me under my circumstance*, then there is no room for one to represent things that do not possess such a property. Even if dispositionalist Russellian representationalism coheres with virtual veridicalism in this way, the account is trivial because it predicts that no VR experience is falsidical. This consequence is unacceptable to both dispositionalist Russellian representationalism and virtual veridicalism. Unless there are resources to resolve the impossibility of misrepresentation, it is unacceptable for dispositionalist Russellian representationalism to take the indexical strategy to characterize both "subject groups" and "perceptual circumstances" simultaneously.

[10] To forecast again, the account I will propose in Chapter 5 involves the representation of self-locating contents and centering features.

To summarize this section, I conclude that dispositional Russellian representationalism is inconsistent with virtual veridicalism. In brief, if perceptual experience represents dispositional properties, it must specify how the dispositional properties are individuated with respect to both "subject groups" and "perceptual circumstances", and no possible way of individuating "perceptual circumstances" renders this account consistent with virtual veridicalism. Therefore, dispositional Russellian representationalism could not help support virtual veridicalism.

4.4. Fregean Representationalism

Having rejected physical and dispositionalist Russellian representationalism, this section assesses whether Fregean representationalism coheres with virtual veridicalism. I argue that although Fregean representationalism satisfies the constraints 1 – 4 listed in Section 4.2, it fails to satisfy the fifth constraint, *viz. Scope*. As a result, Fregean representationalism is inconsistent with virtual veridicalism.

4.4.1. Fregean Phenomenal Content

While committing to *Representationalism*, Fregean representationalism rejects *Russellianism*. Instead, this view takes phenomenal contents to be Fregean propositions, i.e., propositions composed of *modes of presentation* (MoPs) of their extensions. In other words, for Fregean representationalism, there are two levels of phenomenal content: MoPs and extensions. Before going to Fregean representationalism, let me clarify what are MoPs.

The notion of MoP is derived from Frege (1892), where he tried to resolve a puzzle regarding linguistic opacity. Consider his famous example: The name "Phosphorus" refers to the first star to appear in the morning sky, and the name "Hesperus" refers to the first star to appear in the evening sky. In fact, both names refer to the planet Venus. Traditionally, the content of a name is nothing over and above its extension. So, it follows that substituting "Phosphorus" for "Hesperus" in a statement always preserves the truth value of the statement, since they mean the same object, Venus. However, this is not the case. Although "Hesperus is Hesperus" is necessarily true, "Hesperus is Phosphorus" is not so since, intuitively, Hesperus could not have been Phosphorus. As Frege analyzes, the difference in our intuition about these statements is due to the difference in their cognitive significance in content. He introduces the notion of MoPs to capture this element of content that is in addition to extension. Presumably, "Hesperus is Hesperus" is not informative

or cognitively significant, but "Hesperus is Phosphorus" is cognitively significant. According to Frege, this is because "Hesperus" and "Phosphorus" in the second sentence have different MoPs, and the difference in their MoPs determines the difference in their contents despite the same extension.

There are different ways to think about the relationship between MoPs and extension under the broadly Fregean framework. A typical way of understanding Fregean MoPs is that MoPs determine extensions. This entails that two representations with the same MoP must have the same extension. The two-dimensional semantic theories (e.g., Chalmers, 1996, 2002, 2006), however, provide another way to characterize the notion. MoPs can be thought of as primary intensions, which are functions from worlds considered as actual to extensions. The same function in different centered worlds can return different extensions. For example, the primary intension of "water" picks out as extension H_2O in the actual world but XYZ on Twin-Earth. This reflects the intuition that if it turns out that Twin-Earth is the actual world, it will turn out that "water" refers to XYZ. In this way, two representations may share the same MoP but have different extensions, depending on the worlds they are situated.

The two-dimensionalist MoPs of linguistic content can be extended to characterize phenomenal content. Different versions of Fregean representationalism differ in their characterization of MoPs of experiences (i.e., phenomenal MoPs). In the rest of this chapter, I focus on Chalmers's (2004, 2006) and Thompson's (2009) version because this version is the most developed one in the literature. So, unless specified otherwise, I reserve the label "Fregean representationalism" for the particular Chalmers-Thompson view. According to this view, phenomenal MoPs are characterized as *conditions on extension* (CoEs), i.e., conditions that an object or property must satisfy in order for it to be the extension of that experience. Specifically, an experience involving certain phenomenal content has a phenomenal MoP that imposes CoE like *the property that normally causes experiences of that sort in me*. In this view, normality is statistical typicality. Statistical typicality is a backward-looking notion, requiring accumulating many past tokening instances of a certain experience type (Thompson, 2009). For instance, C_{700} is the statistically typical cause of my reddish experience because most of my past token reddish experiences were caused by C_{700}. Like Egan (2006a, 2006b), Fregean representationalism also employs an indexical strategy to the "in me" part involved in phenomenal MoPs. They both appeal to the notion of centered worlds, but they have different views on the notion. Fregean

representationalists endorse two-dimensional semantics, while Egan does not. According to two-dimensional semantics, an indexical element is characterized by its primary intension (Chalmers, 2002). The primary intension of "I" can be understood as picking out the individual marked at the center of any given possible world. Centered on David Chalmers, "I" returns David Chalmers as its referent. The same primary intension picks out Brad Thompson if centered on Brad Thompson. In this sense, the primary intension of "I" remains the same across different expressions involving "I". Therefore, a phenomenal MoP is characterized as "a function from a centred world that picks out the property that typically causes [the type of] experiences in the individual at that centre" (Thompson, 2009: 108; see also Chalmers, 2004).

Recall Jack and Jill, for example, the phenomenal MoP of Jack's reddish experience picks out C_{700} as extension because C_{700} is the typical cause of reddish experience in the individual marked at the center, namely Jack. In contrast, the phenomenal MoP of Jill's reddish experience picks out C_{550} as extension because C_{550} is the typical cause of reddish experience in the individual marked at the center, namely Jill. Hence, although Jack and Jill share the same phenomenal MoPs (i.e., phenomenal redness), their MoPs pick out different properties as extensions.

Although characterizing phenomenal MoPs as imposing CoEs like *the property that normally causes experiences of a certain sort in me* resolves cases like spectrum inversion, it faces an immediate difficulty while explaining cases like color constancy (Millar, 2013: 223). Due to color constancy, one's experiences of different objects under different lighting conditions can share the same phenomenal character but represent them as having different colors. Suppose that I have phenomenally identical yellowish experiences while looking at a yellow sheet of paper under the fluorescent light and while looking at a white sheet of paper under the yellow light. Throughout my perceptual history, it is plausible that what typically causes my yellowish experience is the yellow property rather than the white property. If the MoP of my yellowish experiences picks out the property that typically causes yellowish experiences in me, then I must misrepresent the white paper as having the yellow property even under the yellow light.

Thus, to accommodate color constancy, Fregean representationalists need to add something to their account of phenomenal MoPs. Thompson (2009: 110–2) argues that the difficulty cannot be resolved simply by building in perceptual conditions to CoE, such that phenomenal MoPs pick out something like *the property that normally causes yellowish experiences in me under a certain lighting condition*. If so, the MoP of my yellowish experience under the fluorescent light picks out

the property that normally causes yellowish experiences in me under the fluorescent condition, whereas that of my yellowish experience under the yellow light picks out *the property that normally causes yellowish experiences in me under the yellow light condition*: They do not share the same MoP. Nor can Fregean representationalists resolve the difficulty by specifying lighting conditions in an indexical way, such that phenomenal MoPs pick out something like *the property that normally causes yellowish experiences in me under the current lighting condition*, since this would make misrepresentation impossible.[11]

Instead, Thompson appeals to a holistic conception of phenomenal content, arguing that "phenomenal modes of presentation for colour experiences might be sensitive to properties of the entire visual field in their determination of reference" (2009: 112). In this view, the MoP of my yellowish experience picks out the property that typically causes yellowish experiences in me in instances where the total phenomenology of the yellowish experiences is similar to the total phenomenology of the present experience. When I look at the yellow paper under the fluorescent light, the surroundings will presumably have the appearance that is typical of the surroundings under the fluorescent light. By contrast, when I look at the white paper under the yellow light, the surroundings will presumably have the yellowish appearance that is typical of the surroundings under the yellow light. So, the difference in the total phenomenology indicates that the two yellowish experiences occur in different lighting conditions and, in turn, pick out different properties as extension despite the same MoP. Thompson further argues that this solution can avoid the objection from impossible misrepresentation. If my yellowish experience of the white paper and that of the yellow paper share the same total phenomenology, then the MoP they share picks out the very same property, given that particular total phenomenology. Hence, at least one experience misrepresents. In other words, Thompson commits to the idea that "any difference in lighting conditions that makes a difference to the colour properties [my] experience attributes to objects will be accompanied by some difference in the total phenomenology of [my] experience" (Millar, 2013: 224).

Now, we are ready to see why Fregean representationalism can accommodate the first four constraints of virtual veridicalism listed in Section 4.1. Recall the E_P/E_V pair. E_P comes from a group of one's token reddish experiences that occurs in a situation where she is in an ordinary,

[11] For details, please see Thompson (2009: 110–111). The objection from impossible misrepresentation can be constructed in the way I did in Section 4.3.1.

physical environment, and this group of experiences is reliably caused by C_{700}; E_V comes from another group of her token reddish experiences that occurs in a situation where she is regularly immersed in a virtual world, and that group of experiences is reliably caused by VC_{700}. Since current VR technology is unable to simulate a scenario that appears completely the same as the ordinary world, I find no reason to deny that one's reddish experiences in an ordinary environment and those in a virtual world may have different total phenomenology.[12] Given the difference in the total phenomenology of E_P and E_V, according to Thompson's view, they occur under different conditions where the phenomenal MoP of reddish experience picks out different properties as extension. With this respect, E_P and E_V pick out different extensions but still share the same phenomenal MoP because they are both reddish experiences. *Same Phenomenology* is satisfied. Secondly, the properties that E_P and E_V represent are the normal cause of reddish experience. C_{700} satisfies the CoE imposed by the MoP of reddish experience in the ordinary environment because C_{700} normally causes reddish experience in the ordinary environment. VC_{700} satisfies the CoE imposed by the MoP of reddish experience in the virtual world because VC_{700} normally causes reddish experience in the virtual world. Hence, *Different Reliable Cause* is also satisfied. Thirdly, the veridicality condition is fixed by the extension. E_P is veridical in the ordinary environment iff its object instantiates C_{700}, and E_V is veridical in the virtual world iff its object instantiates VC_{700}. Since C_{700} and VC_{700} are instantiated in the ordinary and virtual environment, respectively, E_P and E_V are veridical under their respective circumstances. *Veridical non-VR Experience* and *Veridical VR Experience* are both satisfied. Therefore, Fregean representationalism satisfies the first four constraints of virtual veridicalism, cohering with virtual veridicalism in perfect virtual worlds.

4.4.2. The Problem of Mixed Reality

Although Fregean representationalism satisfies the first four constraints of virtual veridicalism, it is unclear whether this account can also satisfy the fifth constraint, *viz. Scope*. If Fregean representationalism fails to satisfy *Scope*, it fails to be an appropriate theory of content that coheres with the version of virtual veridicalism defended in Chapter 3. This sub-section argues that Fregean representationalism cannot satisfy this constraint because this account does not contain resources to accommodate veridical perception in mixed worlds.

[12] For instance, Chalmers (2017: 331) suggests that our VR experiences will contain the *phenomenology of virtuality*.

According to Fregean representationalism, the CoE of an experience is satisfied by the property that normally causes that experience, and the property may differ depending on what the environment turns out to be. Since the satisfaction of an experience's CoE determines its veridicality condition, the latter cannot be determined if the former is undetermined. If the extension of phenomenal redness is undetermined, there is no hope of determining the veridicality of reddish experience.

Fregean representationalism predicts veridical perception in a perfect virtual world, where all the perceived objects and properties are virtual. In a perfect virtual world, the CoE of phenomenal redness is clear-cut: It is VC_{700}, not anything else, that normally causes reddish experiences. So, phenomenal redness picks out VC_{700} as extension in a perfect virtual world, and a reddish experience is veridical in a perfect virtual world iff its object instantiates VC_{700}. However, as I am proceeding to argue, such a prediction does not hold in a mixed world, where both virtual and non-virtual properties exist and play a significant role in forming perceptual experiences, since Fregean representationalism implies that no perceptual experience is veridical in mixed worlds.

Recall Sally, the sophisticated VR user addressed in Section 3.5. She wears advanced glasses that cover her entire visual field for a long time. Ordinarily, the glasses function as common glasses. When Sally wants to enter VR, she switches on the glasses. In doing so, the lenses block the array of lights from the outside, and the simulated scenario is shown on the lens. Now, Sally starts perceiving virtual things in a V-100% scenario[13] because the simulation completely screens off non-virtual entities through the lens. Moreover, the VR glasses are so advanced that they make a perfect simulation that appears exactly the same as the appearance without wearing the glasses. Whenever Sally switches on the glasses, the appearance of the upcoming V-100% scenario is identical to that of the previous V-0% scenario, so she cannot distinguish the V-100% scenario from the V-0% scenario just by their visual appearances.

Suppose that at a particular moment, what appears to Sally involves two phenomenally identical red circle patches, one on the right (call it ***R***) and the other on the left (call it ***L***). In the V-0% scenario, i.e., before the glasses are switched on, the objects causing the appearance of ***R*** and ***L*** are physical objects, and the property they instantiate is C_{700} (see **Figure 4-1A**). In contrast,

[13] To refresh our memory, I refer to a scenario with a certain degree of Reality-Virtuality mixture by its virtuality percentage. For instance, the V-0% scenario is a perfectly non-virtual environment (a.k.a. an ordinary environment), and the V-100% scenario is a perfect virtual world. All scenarios from the V-1% scenario to the V-99% scenario are mixed worlds. For details, please see Section 3.5.

in the V-100% scenario, i.e., after the glasses are switched on, the objects causing the appearance of ***R*** and ***L*** are virtual objects, and the property they instantiate is VC_{700} (see **Figure 4-1B**). In short, the pair of objects causing the appearance of ***R*** and ***L*** in Sally's experience in the V-0% scenario is different from the pair of objects in the V-100% scenario. Given the plausible assumption that Sally's experiences of ***R*** and ***L*** (call them ***R***-***L*** experiences) in the V-100% scenario and in the V-0% scenario differ in their total phenomenology, it suffices to grant that, according to Thompson's view, the property that normally causes Sally's ***R***-***L*** experience differs in these scenarios. So, Sally's ***R***-***L*** experience can be veridical in both scenarios. On the one hand, the ***R***-***L*** experience represents C_{700} in the V-0% scenario because C_{700} is the normal cause of phenomenal redness in the V-0% scenario. On the other hand, it represents VC_{700} in the V-100% scenario because VC_{700} is the normal cause of phenomenal redness in the V-100% scenario. Virtual veridicalism is well supported by the Fregean explanation. So far, so good.

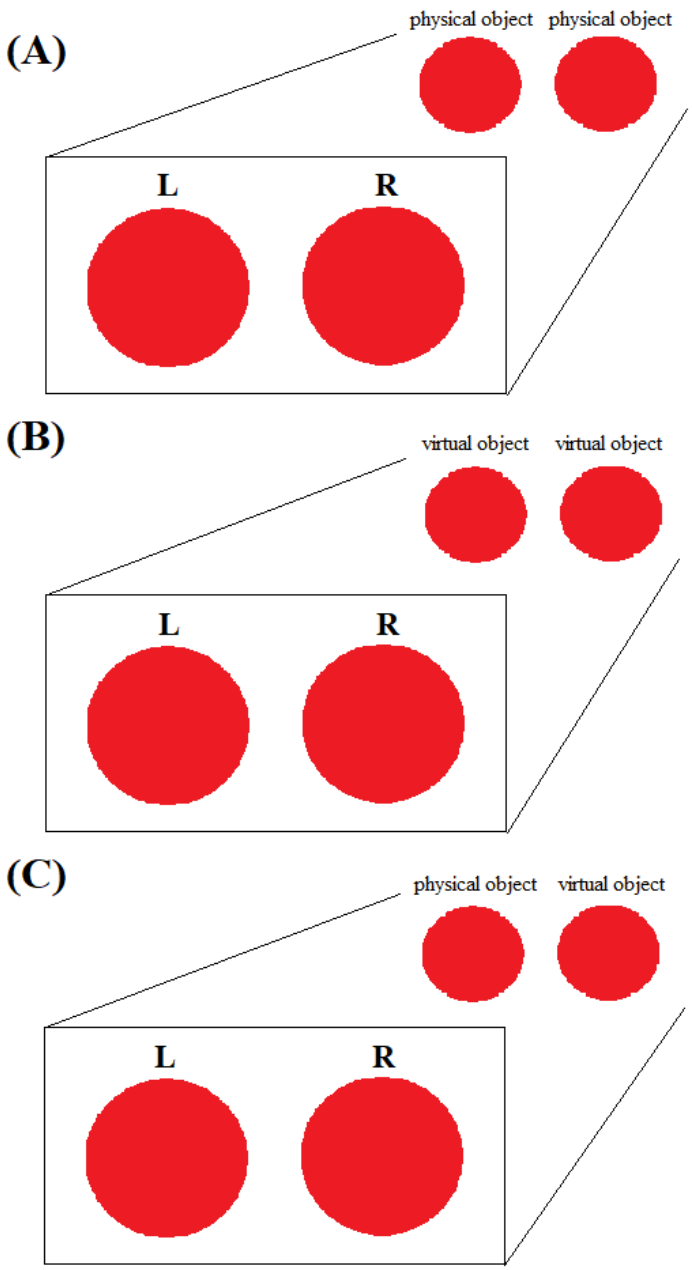

Figure 4-1. Illustrations of the appearance of Sally's ***R***-***L*** experience. (The rectangle box indicates Sally's visual field.) (A) V-0% scenario: Before the VR function is switched on, the objects causing the appearance of ***R*** and ***L*** are physical objects, and the property they instantiate is C_{700}. (B) V-100% scenario: After the VR function is switched on, the objects causing the appearance of ***R*** and ***L*** are virtual objects, and the property they instantiate is VC_{700}. (C) Mixed world: After the MR function is switched on, the object causing ***R***'s appearance is a virtual object, and the property it instantiates is VC_{700}. Yet, the object causing ***L***'s appearance remains a physical object, and the property it instantiates is C_{700}.

However, a problem arises when the story continues. Not only can Sally's glasses simulate a V-100% scenario that shuts out any information from the non-virtual environment, but they also make a partial simulation that blends virtual things into the non-virtual environment, hence generating a mixed world. This is the case of mixed reality (MR). When the MR function is switched on, the glasses block a significant proportion of the non-virtual environment and replace it with virtual counterparts. In the case of Sally's ***R***-***L*** experience, the glasses replace the object that causes ***R*** with a virtual counterpart while leaving ***L*** unchanged (see **Figure 4-1C**). So, for Sally's ***R***-***L*** experience in the mixed world, the object causing ***R*** is a virtual object, and the property that causes phenomenal redness in ***R*** is VC_{700}. In contrast, the object causing ***L*** is still a physical object instantiating C_{700}. This is a case in which a single experience contains two local regions of its quality space that share identical phenomenal character caused by different external properties.

Now, the problem is that the ***R***-***L*** experience cannot be veridical in any mixed world. Do the properties that cause local phenomenal redness in each area satisfy the CoEs imposed by the phenomenal redness? Note that the ***R***-***L*** experience is a single experience in a single subject. Since phenomenal redness picks out as extension the property that normally causes reddish experience, it cannot simultaneously pick out as extension the property that is not the normal cause. If C_{700} normally causes reddish experience, then ***R*** and ***L*** pick out C_{700} as extension. As the property that causes ***R*** is VC_{700}, the ***R***-***L*** experience must be falsidical. Similarly, if VC_{700} normally causes reddish experience, then ***R*** and ***L*** simultaneously pick out VC_{700} as extension. As the property that causes ***L*** is C_{700}, the ***R***-***L*** experience is still falsidical. This argument is generalized as follows:

MR1 The ***R***-***L*** experience is veridical only if both ***R*** and ***L*** are veridical.
MR2 The property that causes ***R*** differs from the property that causes ***L***.
MR3 In the ***R***-***L*** experience, ***R*** and ***L*** pick out the same property as extension.
MR4 Therefore, the ***R***-***L*** experience cannot be veridical.
MR5 If **MR4** is true, most experiences in any mixed world are falsidical.
MR6 Therefore, most experiences in any mixed world are falsidical.

MR1 is implied by the common assumption of what a veridical experience should meet. It is widely accepted that an experience is veridical only if everything involved in its content is veridical. For example, in a paradigmatic illusion where a stick appears bent when partly immersed

in the water, what is mistaken in our experience is only the property of *being bent*. Many other properties are veridically represented, such as *being brown*, *being made of wood*, and so on. Even so, we consider the experience falsidical because something involved in its content is not veridical. Likewise, if either ***R*** or ***L*** misrepresents, we will consider the ***R***-***L*** experience falsidical.

MR2 is the setting of Sally's story. The mixed world in the story is not only conceivable and metaphysically possible but also nomologically possible. I do not think there is a good reason to reject it, as the current VR technology has signaled such a future possibility.[14]

MR3 is the consequence of Fregean representationalism (i.e., the Chalmers-Thompson view). As discussed above, phenomenal redness picks out as extension the property that normally causes reddish experience. Since ***R*** and ***L*** share the same phenomenal MoP, i.e., phenomenal redness, they also pick out as extension the property that normally causes reddish experience. So, if one accepts Fregean representationalism, one should accept MR3.

There are potential objections to MR3. First, a Fregean representationalist might argue that in the mixed world, both C_{700} and VC_{700} are the normal cause of phenomenal redness because they both count as statistically typical causes of reddish experiences. If so, the extension of phenomenal redness in the mixed world is a *disjunction* in the form of C_{700} *or* VC_{700}. True, normality as well as statistic typicality are not exclusive notions. Although this response is possible, it owes us an explanation of *why* two such heterogeneous properties constitute a common content, as it would not be helpful to say that, *de facto*, they cause the same experience type in the mixed world.[15] Thus, the response from disjunctive content seems a suboptimal solution.

Second, a Fregean representationalist might reject MR3 by arguing that phenomenal MoPs involve demonstrative elements that enable ***R*** and ***L*** to pick out different properties. For instance, ***R*** may pick out the property that *normally causes reddish experiences at THIS region of quality space in me*, while ***L*** may pick out the property that *normally causes reddish experiences at THAT region of quality space in me*. If so, ***R*** can pick out VC_{700} while ***L*** still picks out C_{700}. However, this objection leads to counterintuitive consequences. First, it is simply unclear how the phenomenal content involves demonstrative elements. Second, the individuation of a phenomenal character would become extremely fine-grained. In particular, the reddish character would no

[14] In fact, VR users have been able to interact with virtual objects that are mixed into non-virtual environments through goggles like the *Microsoft HoloLens 2*. I believe that the mixed world that yields the ***R***-***L*** experience can also be easily produced.

[15] I will address the demand of such explanation in more detail in Section 5.4.1.

longer be a character type but merely a determinable of distinct determinates such as reddish-at-*THIS-region* and reddish-at-*THAT-region*. This is quite counterintuitive because colors *per se* do not involve spatial elements as constituents. Hence, the objection from demonstrative elements involved in Fregean content does not help solve the problem.

Third, one might take a weaker claim than the initial formulation and argue that phenomenal redness picks out *whatever property normally causes reddish experience in me*. In this sense, the property picked out by the same MoP need not be unique. The CoE can be simultaneously satisfied by both C_{700} and VC_{700} so that the ***R***-***L*** experience can be veridical. However, this is not the case, either. The CoE can be satisfied depending on different readings of the "property" in question: *Whatever property normally causes reddish experience in me* can be read *de dicto* or *de re*. If it is read *de dicto*, the phenomenal MoP no longer fixes the extension on particular properties but on their description. This suggests that a third level of content other than MoPs and extensions is introduced to characterize phenomenal content. This is unacceptable to Fregean representationalists, who presumably would not give up the Fregean two-level content framework.[16]

So, *whatever property normally causes reddish experience* should be read *de re* to avoid violating the Fregean framework. But the *de re* reading does not help make the CoE simultaneously satisfied by both C_{700} and VC_{700}. If reddish experience represents *whatever property normally causes reddish experience in me* (*de re*), then the CoE has to pick out the specific property that normally causes reddish experience. It could be either C_{700} or VC_{700}, but not both. Either way, the ***R***-***L*** experience might be veridical either when the properties that cause ***R*** and ***L*** are both C_{700} or when the properties that cause ***R*** and ***L*** are both VC_{700}. However, since the CoE cannot pick out C_{700} and VC_{700} simultaneously, the ***R***-***L*** experience cannot be veridical when the property that causes ***R*** is C_{700} while the property that causes ***L*** is VC_{700}, or when the property that causes ***R*** is VC_{700} while the property that causes ***L*** is C_{700}. Therefore, the ***R***-***L*** experience is still predicted to be falsidical. This response does not save Fregean representationalism, either.

[16] If a contender does not stick to the Fregean framework, she will be happy to accept the "whatever property" strategy. This will help her reject MR3. I am also happy with this consequence because this section aims to criticize Fregean representationalists, namely those who follow the Fregean framework. In fact, the "whatever property" strategy resembles what I label as realizer representationalism, according to which phenomenal content attributes the property that plays a role of normally causing experiences of a certain sort in the subject under current perceptual circumstance. I will address and raise an objection to this view in Section 5.4.1.

Due to the truth of MR1 – MR3, Fregean representationalists must accept MR4: The ***R-L*** experience cannot be veridical. If MR5 is true, MR6 follows, then Fregean representationalists should accept that most experiences in any mixed world are falsidical. Consequently, virtual veridicalism does not hold under the Fregean representationalism in any mixed world.

Now, I turn to justify MR5: If MR4 is true, most experiences in any mixed world are falsidical. Fregean representationalism claims that phenomenal redness picks out as extension the property that normally causes reddish experience. In the mixed worlds with different degrees of Reality-Virtuality mixture, the normal cause of reddish experience differs. Recall Milgram and Kishino's (1994) Reality-Virtuality (RV) continuum discussed in Section 3.5. There are virtual worlds including the V-0% scenario, the V-1% scenario, …, and the V-100% scenario. Provided that C_{700} outnumbers VC_{700} in generating reddish experiences in scenarios from the V-0% scenario to the V-49% scenario, C_{700} is the statistically typical cause of reddish experiences in these scenarios. According to Thompson's view of normality, C_{700} counts as the normal cause of reddish experience in the subject in these scenarios, so the extension of the phenomenal redness in these scenarios is C_{700}. Similarly, from the V-51% scenario to the V-100% scenario, the extension of phenomenal redness is VC_{700}. Yet, the extension of phenomenal redness in the V-50% world is indeterminate because which property counts as the normal cause is indeterminate. Thus, we get **Figure 4-2**.

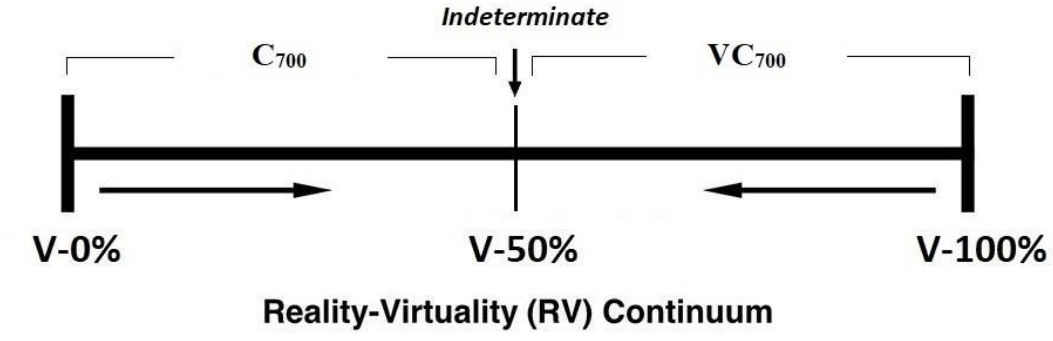

Figure 4-2. The extension of phenomenal redness within the Reality-Virtuality Continuum.

For ease of discussion, imagine an experience containing one hundred phenomenally identical red circle patches ordered in a row (call it the ***100-patches*** experience). From left to right, call the first patch $\boldsymbol{R_1}$, the second patch $\boldsymbol{R_2}$, …, and the last patch $\boldsymbol{R_{100}}$. The ***100-patches*** experience

is generated in every scenario within the RV continuum. In the V-0% scenario, all patches are caused by C_{700}. In the V-1% scenario, only $\boldsymbol{R_1}$ is caused by VC_{700}. In the V-2% scenario, only $\boldsymbol{R_1}$ and $\boldsymbol{R_2}$ are caused by VC_{700}. In the V-3% scenario, only $\boldsymbol{R_1}$, $\boldsymbol{R_2}$, and $\boldsymbol{R_3}$ are caused by VC_{700}. And so on. In the V-100% scenario, all patches from $\boldsymbol{R_1}$ to $\boldsymbol{R_{100}}$ are caused by VC_{700}.

The veridicality of the ***100-patches*** experience within this continuum is determined in terms of the extension in each scenario (see **Figure 4-4**). In the V-0% scenario, the extension of phenomenal redness is C_{700}, and all red patches in this scenario are caused by C_{700}, so the ***100-patches*** experience is veridical. Likewise, the ***100-patches*** experience in the V-100% scenario is also veridical, since all red patches are caused by the property that satisfies as the extension of phenomenal redness, *viz.* VC_{700}. However, the ***100-patches*** experience cannot be veridical in any mixed world (i.e., scenarios from the V-1% scenario to the V-99% scenario). From the V-1% scenario to the V-49% scenario, the extension of phenomenal redness is C_{700}, but a certain amount of red patches is caused by VC_{700}, so there is always a mismatch between the extension and the cause. A similar mismatch occurs in the scenarios where the extension of phenomenal redness is VC_{700} (i.e., from the V-51% scenario to the V-99% scenario), as a certain amount of red patches is caused by C_{700}. Worst, in the V-50% scenario, the veridicality is undefined because the extension is indeterminate. Overall, the ***100-patches*** experience is veridical only in the V-0% and the V-100% scenarios.

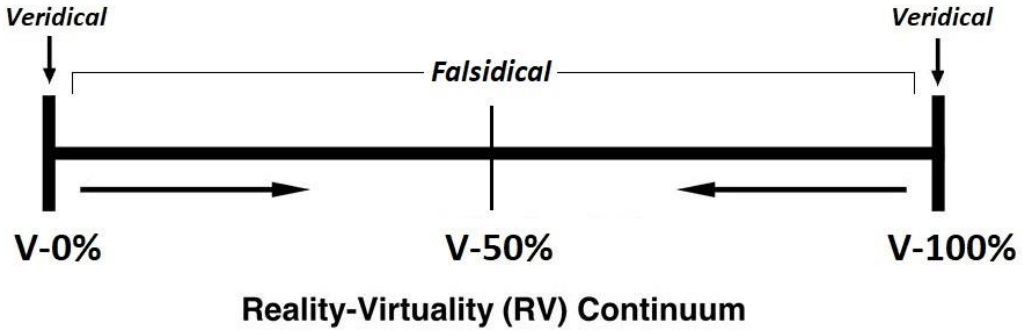

Figure 4-3. The veridicality of the ***100-patches*** experience within the Reality-Virtuality Continuum.[17]

[17] Strictly speaking, a false truth-value is different from an undefined truth-value. But this difference does not matter to my argument. I ignore the difference and take the ***100-patches*** experience in the V-50% scenario as a falsidical experience.

Therefore, according to Fregean representationalism, the ***100-patches*** experience is veridical only in the V-0% scenario and the V-100% scenario. The fact that the ***100-patches*** experience is falsidical in scenarios from the V-1% scenario to the V-99% scenario can be generalized to other experiences containing phenomenally identical local contents caused by distinct properties. MR5 is plausibly true. Then MR6 follows. This suggests that Fregean representationalism does not satisfy *Scope* because *Veridical VR Experience* is not satisfied in any mixed world. Eventually, I conclude that Fregean representationalism is still inconsistent with virtual veridicalism as I defended in Chapter 3.

4.5. Conclusion

This chapter examined the coherence of three currently prominent representationalist accounts and the version of virtual veridicalism defended in Chapter 3. I conclude that none of them is coherent with the view. Physical Russellian representationalism fails because it sticks to physical properties as the representata of perceptual experiences. Dispositionalist Russellian representationalism cannot stably accommodate the constraints of virtual veridicalism because of its irresolvable internal tension of relativization of perceptual circumstances. Fregean representationalism avoids the problems faced by the previous two accounts. However, it does not contain resources to justify the veridicality of VR experience in mixed worlds. The next chapter will propose a new representationalist account and argue that it perfectly resolves the problems encountered by others.

Chapter 5: Role Representationalism About Perceptual Experience

In search of a representationalist picture of lenient content that coheres with the version of virtual veridicalism defended in Chapter 3, Chapter 4 concluded that no currently prominent accounts succeed in providing a satisfying picture. This chapter proposes a new account — *role representationalism* — and argues that this account perfectly coheres with the view.

To reiterate the terminology, perceptual experience has two types of properties. For a perceptual experience to instantiate a *phenomenal property* is for it to exhibit a certain phenomenal character, which is a distinctive what-it-is-likeness for the subject to undergo that experience. For a perceptual experience to instantiate a *representational property* is for it to represent a certain *content.* There are two ways to associate content with an experience: The *strict content* represents the property that appears in the experiences' phenomenology; while the *lenient content* represents the property that is partly determined by the appropriate contingent relationship between the experiences and the external world. I take contents of experiences as their veridicality conditions, i.e., the conditions under which the experiences are veridical. So, there are also two senses with respect to which an experience is veridical: An experience is *strictly veridical* iff its strict content is true, while it is *leniently veridical* iff its lenient content is true. Although this chapter looks for an account of lenient content, it will also address the relation between lenient content and strict content.

Again, my defense of virtual veridicalism focuses on low-level phenomenal characters (e.g., colors, locations, and shapes) and their phenomenal contents (i.e., contents that are determined by phenomenal characters). Here, I assume the *phenomenal content view*, i.e., for any experience with low-level phenomenal characters, there are contents determined by these characters. Apparently, strict contents are phenomenal contents because they attribute the properties that appear in phenomenology. For the first pass, I also stipulate that lenient contents are phenomenal contents since, *prima facie*, experiences of the same sort always play the same experiential role. I will discuss whether they exactly count as phenomenal contents in Section 5.2.5.

Following Chapter 4, this chapter operates within the representationalist framework. *Representationalism* is the view that every phenomenal property is determined by some representational property. This chapter proposes a new account — *role representationalism* — and argues that this account perfectly coheres with the version of virtual veridicalism defended in

Chapter 3. For terminological brevity again, unless specified, I mean this specific version by "virtual veridicalism". Section 5.1 reviews the constraints of virtual veridicalism. Section 5.2 proposes and articulates role representationalism. Section 5.3 details why role representationalism coheres with virtual veridicalism. Section 5.4 addresses potential objections to role representationalism. Eventually, I will conclude that role representationalism and virtual veridicalism are mutually supporting.

5.1. Virtual Veridicalist Constraints Revisited

Following Section 4.1, my aim is to provide an account of the contents of perceptual experiences in virtual worlds that is coherent with virtual veridicalism. This section briefly reviews the constraints of virtual veridicalism. They will be used to assess role representationalism in the following sections.

Chapter 3 defended virtual veridicalism in terms of the phenomenon of perceptual variation without misperception across virtual worlds and ordinary environments. A gist underlying this phenomenon is that many VR experiences and non-VR experiences are leniently veridical:

1. *Veridical VR Experience.* Many VR experiences are leniently veridical.
2. *Veridical Non-VR Experience.* Many non-VR experiences are leniently veridical.

The case of perceptual variation suggests that two groups of token perceptual experiences of the same type can share the same phenomenal characters while being reliably caused by categorically different properties. Previously, I defined VR experience as the *group* of token experiences of a certain type that occur in a virtual world and non-VR experience as the *group* of token experiences of a certain type that occur in an ordinary environment. In a pair of phenomenally identical VR and non-VR experiences (call them *corresponding* experiences), the VR experience is reliably caused by virtual properties, whereas its corresponding non-VR experience is reliably caused by non-virtual, physical properties:

3. *Same Phenomenology.* A VR experience and its corresponding non-VR experience sometimes share some phenomenal characters.

4. *Different Reliable Cause*. The properties that reliably cause a VR experience are different in kind from those that reliably cause its corresponding non-VR experience.

Finally, as the scope of virtual veridicalism covers not only perfect virtual worlds (i.e., virtual worlds that *completely* screen off non-virtual things and then appear the same as non-VR environments) but also mixed worlds (i.e., virtual worlds that *partly* screen off non-virtual things and then appear the same as non-VR environments), the four constraints above have to be satisfied in both worlds:

5. *Scope. Veridical VR Experience*, *Veridical Non-VR Experience*, *Same Phenomenology*, and *Different Reliable Cause* are satisfied in both perfect virtual worlds and mixed worlds.

This chapter assesses the coherence of role representationalism with virtual veridicalism by examining whether it contains sufficient resources to accommodate all these constraints. Before that, let's discuss role representationalism in detail.

5.2. Role Representationalism about Perceptual Experience

This section proposes and articulates role representationalism. Role representationalism is an account of perceptual experiences. It consists of representationalism and *role functionalism about lenient content*. As representationalism is assumed in these chapters, this section pays more attention to role functionalism.

Role representationalism is an account of the lenient content of perceptual experience. The account is neutral to whether the properties that the perceived objects appear to have are instantiated in our world. It does not matter whether strict contents veridically represent or not, which is the core concern of the error theorists (e.g., Boghossian & Velleman, 1989; Mendelovici, 2013; see also Section 3.2). This advantage applies to other perceptual states, including spatial experience or experiences in other perceptual modalities. I raised a concern in Section 3.3 that the space we perceive the external reality as having may not be identical to the actual physical space, leaving room for some phenomenally spatial properties never to be instantiated in the world. However, even if such a possibility comes true, the fact that the strict contents of spatial experience

do not veridically represent does not imply systematic misperception of space with respect to lenient contents. Since the actual world may play a sufficient role in producing our spatial experience, our spatial experiences can still be leniently veridical. Role representationalism aims to preserve the ordinary conception of perceptual veridicality while being compatible with the possibility that our world does not appear to us in the way it really is.

Role representationalism is motivated to accommodate perceptual variation without misperception. Recall my argument in Chapter 3. We usually have orangish experiences while perceiving a certain property, but we may also reliably perceive another property as having the same color. We usually perceive a certain property as having a certain size. Still, we may also reliably perceive another property as having the same size. We usually perceive a certain shape property as round, but we may also reliably perceive another shape property as having the same shape. The perceptual pairs in these cases strongly suggest that phenomenally identical experiences can be reliably generated by distinct external properties.

Take spectrum inversion without illusion as the first case of perceptual variation without misperception. Suppose two groups of perceivers in our world have systematically inverted color experiences when looking at the same objects. Jack and Jill are normal perceivers of each group. Jack perceives ripe tomatoes as phenomenally red, while Jill perceives grass as phenomenally red. Specifically, Jack's reddish experience is caused by C_{700} (i.e., the property of *having a surface reflectance profile that reflects light with a wavelength of 700nm*), whereas Jill's is caused by C_{550} (i.e., the property of *having a surface reflectance profile that reflects light with a wavelength of 550nm*). Presumably, there is no sense in assuming either Jack's or Jill's reddish experiences to be falsidical. In this Jack/Jill case, perception is varied with respect to different subjects.

Consider another case of perceptual variation without misperception, which was addressed in Section 4.2. Consider two reddish experiences in the same subject, E_P and E_V. E_P is a token of her reddish non-VR experience, which is reliably caused by C_{700} in an ordinary environment, while E_V is a token of her reddish VR experience, which is reliably caused by a virtual property (call it VC_{700}) in a virtual world. Although E_P and E_V come from groups of reddish experiences that are reliably caused by different properties, they share the same phenomenal character. Again, there is no sense in assuming either E_P or E_V to be falsidical. In this E_P/E_V case, perception is varied with respect to different perceptual circumstances.

At first glance, these cases seem to conflict with the phenomenal content view. In the Jack/Jill case, the phenomenal content view would predict that Jack's and Jill's reddish experiences share the same phenomenal content because they share the same phenomenal character. However, as a ripe tomato and grass hardly possess the same color, it seems contradicting to say that both Jack and Jill perceive veridically. Similarly, in the E_P/E_V case, the phenomenal content view would predict that E_P and E_V share the same phenomenal content because they share the same phenomenal character. However, as C_{700} and VC_{700} are distinct properties, it also seems contradicting to say that both E_P and E_V are veridical. Unless the phenomenal content view adopts a certain view of the nature of phenomenal content, it seems hard to accommodate both cases.

Role representationalism aims to resolve this conflict by adopting role functionalism about lenient content. Specifically, the latter view is formulated as follows:

> *Role Functionalism about Lenient Content*: The lenient content of a perceptual experience with phenomenal character *C* is a proposition of the form <something is *F*>, where *F* is the role property of *having a realizer property that normally causes experiences with C in the subject under the current perceptual circumstance*.

I explain each part of this account in the following sub-sections.

5.2.1. Role Property

According to Section 2.3.1, a *role property* is a higher-order property whose instantiation is realized by the instantiation of its first-order realizer property. A role property is individuated in terms of the functional role played by its realizers. The role property attributed by lenient contents is the property individuated by the functional role of normally causing the associated experience type in the subject under the current perceptual circumstance.

5.2.2. Subject

Following the indexical strategy discussed in Section 4.3.1, the "*subject*" part is understood as an indexical element, making phenomenal content a sort of *self-locating content* (Egan, 2006a, 2006b; see also Perry, 1979; Lewis, 1979). Unlike contents that are commonly taken as possible-

worlds propositions, self-locating contents are *centered-worlds* propositions. A centered world is a set-theoretic object that can be understood as a <world, <individual, time>> pair. By representing self-locating content, one represents not only what the world is like but also one's location within the world.[1] By representing <a friend is being nearby>, the self-locating content is the set of centered worlds in which a friend is near the center, i.e., it is all and only <*w*, <*i*, *t*>> pairs such that a friend is near *i* at *t* in *w*. In this way, my representation <a friend is being nearby> and your representation <a friend is being nearby> is the same set of centered worlds, i.e., they share the same self-locating content.

Similar to properties, we can describe self-locating contents in terms of "property analogs", which Egan (2006a, 2006b) calls *centering features*. We describe possible-worlds contents in terms of the properties that things are represented as having. We can also describe self-locating contents in terms of the centering features that things are represented as having. In other words, as we typically think of possible-worlds contents as structured propositions made of properties, we can also think of self-locating contents as structured propositions made of centering features. While properties are typically considered as functions (or things that determine functions) from possible worlds to extensions, centering features can be taken as functions (or things that determine functions) from centered worlds to extensions. Given a world and a center, *being nearby* will assign a set of things that are near the center of that world. Given an object, *being nearby* will give us a set of centered worlds in which that object is near the center. So, by representing a friend as *being nearby*, we pick out a set of centered worlds where those centered worlds are the ones in which a friend is near the center or, more specifically, "the ones that, when plugged into the *being nearby* function, deliver an extension that includes [a friend]" (2006b: 109). In particular, by representing something as being nearby, there is a centering feature corresponding to the self-locating content <nearby> that picks out the set of things near the center of the world. Hence, the fact that my representation <a friend is nearby> and your representation <a friend is nearby> share the same self-locating content is equivalent to their representing the same centering feature.

(Note: Adopting the centering feature approach, a role representationalist should no longer call the representata "role properties" and should replace the label with something like "role

[1] Again, the "location" in this sense need not be geographical. Many self-locating contents such as <my hat is too tight>, <Fozzie is staring at me>, and <limes are disposed to cause greenish experience in me under illuminant circumstance> have nothing to do with geographical self-location (Egan, 2006a: 511).

centering features". For convenience, I still stick to the label "role properties" in the rest of the dissertation. Whenever I address role representationalism, by "role properties", I mean "role centering features".[2])

The indexical strategy helps accommodate the Jack/Jill case. By representing something as having the role property of *having a realizer property that normally causes reddish experiences in me under current perceptual circumstance*, the lenient content picks out all and only <*w*, <*i*, *t*>> pairs such that something has the role property of *having a realizer property that normally causes reddish experiences in* i *under current perceptual circumstance* at t *in* w. When Jack sees a ripe tomato as red, he represents the tomato as having the role property of *having a realizer property that normally causes reddish experiences in me under the current perceptual circumstance*. When Jill sees grass as red, she represents the grass as having the role property of *having a realizer property that normally causes reddish experiences in me under the current perceptual circumstance*. On the one hand, the phenomenal content view is satisfied because Jack and Jill share the same lenient content. On the other hand, the intuitive force of perceptual variation without misperception is preserved. Firstly, Jack's and Jill's reddish experiences share the reddish character. Secondly, the same role property attributed by their lenient contents is realized by different properties: Jack's reddish experience is caused by C_{700}, whereas Jill's is caused by C_{550} (i.e., the property of *having a surface reflectance profile that reflects light with a wavelength of 550nm*). Thirdly, Jack's and Jill's reddish experiences are probably veridical. Since Jack and Jill are normal perceivers within their perceptual groups, C_{700} and C_{550} do normally cause Jack's and Jill's reddish experiences, respectively. Thus, there is no reason to deny that the ripe tomato and the grass do not instantiate the role property in this case.

5.2.3. Perceptual Circumstance

A *perceptual circumstance* refers to the relevant factors that influence experiences with a certain phenomenal character in a given environment. These factors differ in their relevance with respect to different kinds of phenomenal characters. For instance, illumination is a relevant factor for color experiences, whereas the distance of objects is a relevant factor for size experiences. Relevant factors roughly decompose into environmental factors, factors about perceivers'

[2] Egan (2006b) extensively argues for the metaphysical legitimacy of centering features. Be he successful, I think there is no problem to characterize a role centering feature in the way we characterize a role property.

perceptual apparatus, and factors about perceivers' perceptual hardwiring that eventually yields perceptual experiences. A difference in a relevant factor results in different perceptual circumstances.

Environmental factors, such as the object's illumination condition and surface reflectance profile, determine the configuration of distal stimuli (i.e., the actual object or property in the environment that stimulates a sense organ) that make a phenomenal difference in experiences of a certain sort, e.g., color experiences. *Ceteris paribus*, objects under sunlight and objects under red light look different in color. Also, *ceteris paribus*, objects differing in surface reflectance profile affect color experiences. Objects with C_{700} and objects with C_{550} reflect light with different wavelengths, and such a difference generates different color experiences.

A difference in perceivers' perceptual apparatus also helps differentiate perceptual circumstances. Perceptual apparatus denotes any instrument or equipment used in perception that determines the causal relations from distal stimuli to proximal stimuli (i.e., the physical energy that directly stimulates a receptor or a sense organ). Mirrors, microscopes, and VR headsets are sorts of perceptual apparatus. Suppose one sees things with the aid of myopia glasses. *Ceteris paribus*, the light energy reflected by an object that stimulates the photoreceptors of her retina differs with and without mirrors, affecting the spatial configuration of the object in experience. So, seeing objects with and without mirrors counts as different perceptual circumstances. Microscopes also count as a perceptual circumstance by bringing nearly unobservable properties into our visual experiences. For example, a silky hair may look split under the microscope. Likewise, VR headsets also change the way the distal stimuli cause the proximal ones. Suppose that a virtual property looks red through VR headsets. Without the equipment, the property probably may not look red (or even unperceivable). Still, it is plausible to say that, *ceteris paribus*, the proximal stimuli seen through VR headsets are different from those seen without them. Using VR headsets counts as a relevant factor of a perceptual circumstance.

The third category of relevant factors concerns factors about perceivers' perceptual hardwiring. By perceptual hardwiring, I mean the causal process in the subject's head, starting from the stimulation of a receptor or sense organ to the instantiation of relevant experiences. It is possible for two subjects to experience differently while having all the same stimulation on their receptors or sense organs; it is also possible for them to experience exactly the same while having different stimulations. Paradigmatic cases of this sort are the *Three Planets* (*Color-inverted Earth*,

Doubled Earth, and *El Greco Earth*), which were extensively discussed in Chapter 3. Recall *Color-inverted Earth* in Section 3.2.[3] A gist of this case is that Earthlings and Color-inverted Earthlings are color-inverted. This suggests that, given the same stimulation on their color receptors, they will instantiate inverted color experiences. Due to this difference, we can say that the difference between Earthlings' and Color-inverted Earthlings' perceptual hardwiring is a relevant factor for color experiences. Thus, Earth and Color-inverted Earth are different perceptual circumstances for color experiences. Similarly, in *Doubled Earth* and *El Greco Earth*, the difference between Earthlings' and Doubled Earthlings' perceptual hardwiring and the difference between Earthlings' and El Greco Earthlings' perceptual hardwiring are relevant factors for size experiences and shape experiences, respectively.[4] Consequently, Earth and Doubled Earth are different perceptual circumstances for size experiences, and Earth and El Greco Earth are different perceptual circumstances for shape experiences.

Although the abovementioned factors are apparently relevant to determining a perceptual circumstance, there is no clear-cut criterion for determining whether a factor is relevant or not. In fact, which factors count as relevant is an empirical question. Even so, it does not follow that the determination of relevant factors is fully arbitrary. As far as I can see, there are at least two constraints for the relevancy. One constraint is inclusive; the other is exclusive. The inclusive constraint has been extensively discussed: Factors that, if varied from what they actually are while holding everything else as it actually is, make a difference in experiences of a certain sort in a given environment are relevant factors of that sort. Illumination with different colors of light is a

[3] To refresh our memory, I briefly summarize *Color-inverted Earth* in this footnote. Color-inverted Earth is a far-off planet where the color of all things is systematically inverted with respect to the things on Earth. On Color-inverted Earth, basketballs have the surface reflectance profile that reliably causes Earthling's turquoise-bluish experiences, say, the surface reflectance profile that reflects light with a wavelength of 500nm. By contrast, the surface reflectance profile that reliably causes Earthling's orangish experiences reflects light with a wavelength of 600nm. In addition, color-inverted Earthlings perceive colors in the way that is systematically inverted with respect to Earthlings' color perception, such that what is perceived as turquoise-blue in Earthlings is perceived as orange in Color-inverted Earthlings. Thus, while seeing a basketball on their respective planet, Color-inverted Earthlings and Earthlings both perceive it as orange. They share the same phenomenal character in their experiences of basketball, but the experiences are reliably generated by different surface reflectance profiles. Presumably, there is no good reason to say that either Earthlings or Color-inverted Earthlings suffer from systematic misperception.

[4] Memory refreshment again. Doubled Earth is a far-off planet where the size of all things is systematically doubled with respect to the things on Earth. On Doubled Earth, basketballs have the size that are twice as big as the standard ones on Earth. In addition, Doubled Earthlings perceive size in the way that is systematically halved with respect to Earthlings' size perception. Thus, while seeing a basketball on their respective planet, Doubled Earthlings and Earthlings both perceive it as having the same size. They share the same phenomenal character in their experiences of basketball, but the experiences are reliably generated by different size property. Presumably, there is no good reason to say that either Earthlings or Doubled Earthlings suffer from systematic misperception.

factor relevant to color experiences because a change in illumination affects one's color experiences of objects. The distance of objects influences the size one perceives them as having, so it is a relevant factor for size experiences. Likewise, wearing colored lenses results in different color experiences of objects from not wearing them (or wearing transparent ones), so wearing colored lenses is a relevant factor for color experiences. In contrast, the exclusive constraint helps exclude irrelevant factors: Factors that, if varied from what they actually are while holding everything else as it actually is, cannot make a difference in experiences of a certain sort in a given environment are irrelevant factors of that sort. Illumination with different colors of light is an irrelevant factor of size experiences because no change in the color of light can affect how one perceives the size of objects. [5] The shape of objects cannot influence how one perceives their color, so it is an irrelevant factor for color experiences. Similarly, using a mirror cannot give rise to different color experiences of objects from not using it. Thus, a mirror is an irrelevant factor for color experiences, too.

Ideally, a perceptual circumstance can be defined as an n-tuple in the form of <factor 1, factor 2, …, factor n> that includes all relevant factors as its elements. In a very simplified way, a perceptual circumstance for color experience is an n-tuple <surface reflectance profile, illumination with different colors of light, colored lenses, VR headset, Earthlings' or Color-inverted Earthlings' perceptual hardwiring>, and a perceptual circumstance for size experience is an n-tuple <object size, object distance, myopia glasses, microscope, VR headset, Earthlings' or Doubled Earthlings' perceptual hardwiring>. Any difference in the value assigned to any element within the n-tuple results in different perceptual circumstances. For example, <C_{600}, sunlight, no color lenses, no VR headset, Earthlings' perceptual hardwiring> is a perceptual circumstance for color experiences, and <C_{500}, sunlight, blue color lenses, no VR headset, Earthlings' perceptual hardwiring> is another. As for size experiences, <α size, short distance, no myopia glasses, with VR headset, Earthlings' perceptual hardwiring> is a perceptual circumstance, while <α size, long distance, no myopia glasses, no VR headset, Earthlings' perceptual hardwiring> is another.

Role representationalism can also adopt the indexical strategy to characterize the "current perceptual circumstance". As I have argued that lenient contents are self-locating contents, the

[5] One might wonder what if there is no illumination. Surely one cannot see anything without illumination, but the presence/absence of illumination is not a difference maker but merely an enabler of perception of the *size* of objects. The presence/absence of illumination affects whether one perceives an object' size, but it does not affect what size one perceives the object as having. Thus, the presence/absence of illumination is an irrelevant factor of *size* perception.

notion of the "current perceptual circumstance" can be understood as the geographical location where the individual *i* is located at the time *t* in the world *w*. So, by representing something as having the role property of *having a realizer property that normally causes reddish experiences in me under the current perceptual circumstance*, the lenient content picks out all and only <*w*, <*i*, *t*>> pairs such that something has the role property of *having a realizer property that normally causes reddish experiences in* i *under the perceptual circumstance where* i *is located at* t *in* w.

Therefore, the E_P/E_V case can be accommodated in this way. Adopting the indexical strategy, role representationalism predicts that E_P and E_V share the same lenient content because they share the same self-locating component involved in their perceptual circumstances: They both represent their objects as having the role property of *having a realizer property that normally causes reddish experiences in me under the current perceptual circumstance.* The prediction not only satisfies the phenomenal content view in that E_P and E_V share the same phenomenal content but also preserves the intuitive force of perceptual variation without misperception in the E_P/E_V case. Firstly, E_P and E_V share the reddish character. Secondly, they come from groups of experiences that are reliably caused by different properties: E_P is caused by C_{700}, whereas E_V is caused by VC_{700}. Thirdly, E_P and E_V are probably veridical because the role property is instantiated in both ordinary and virtual environments. Specifically, the role property is realized by C_{700} in ordinary environments and VC_{700} in virtual worlds.

5.2.4. Normal Condition

There is more than one way for a perceptual experience to occur under a perceptual circumstance. For example, light with a wavelength of 600 nm may have different origins, the lenses (if any) can be made of different materials, and so on. Let the *perceptual condition* of an experience be the specific way it occurs, where all causally and metaphysically relevant factors, including, say, the origin of the light and the material of the lenses, are counted as part of the way the experience occurs. These causally and metaphysically relevant factors need not be factors that *in fact* make a difference to the experience, but they are nonetheless part of how the experience comes about. For any given perceptual circumstance, some of the conditions occur when the token experience has causes that are typical for that circumstance, and others don't (maybe due to deviant causation or whatever). Call the former the *typical conditions* and the latter *atypical conditions*. In *Color-inverted Earth*, the conditions where orangish experiences are caused by C_{600} are typical

conditions on Earth (i.e., the perceptual circumstance of the n-tuple <C_{600}, sunlight, no color lenses, no VR headset, Earthlings' perceptual hardwiring>), and the conditions where turquoise-bluish experiences are caused by C_{500} are atypical conditions under the same circumstance. On Color-inverted Earth (i.e., the circumstance with the n-tuple <C_{500}, sunlight, no color lenses, no VR headset, Color-inverted Earthlings' perceptual hardwiring>), by contrast, the conditions where orangish experiences are generated are typical conditions, and the conditions where turquoise-bluish experiences are generated are atypical conditions.

Role representationalism requires a notion of normality, not just typicality. Typicality is a descriptive notion, while normality involves normativity. So, role representationalism needs an account of normality to bridge the gap. In the literature, the naturalist accounts reduce normality to naturalized, descriptive notions, while the non-naturalist accounts do not. Among two prominent naturalist accounts, Thompson (2009) reduces normality to a statistical notion, and Millikan (1984) reduces it to a biological notion.

Thompson (2009) offers a naturalist account that claims that normality just is *statistical typicality*. Statistical typicality is a backward-looking notion. A statistically typical condition requires the accumulation of many past tokening instances of a certain experience type under the perceptual circumstance. Accordingly, conditions in which C_{600} causes Earthlings' orangish experience qualify as normal conditions since they are statistically typical conditions where Earthlings' orangish experience is produced on Earth, and they are statistically typical conditions since they are in the majority among all the past tokening instances of Earthlings' orangish experience. Having Thompson's account, role representationalism would predict that Earthlings' orangish experience on Earth represents the role property realized by C_{600} because Earthlings' orangish experience is caused by C_{600} *under normal conditions* on Earth.

As Thompson (2009: 105–6) admits, his naturalist account implies that any experience that is not preceded by a sufficient sequence of experiences in the same circumstance would lack extension. Suppose an Earthling travels to Color-inverted Earth. Her orangish experience represents nothing at the first instances when she perceives things on Color-inverted Earth because the pre-condition for the satisfaction of its extension is not met, i.e., there is *no* statistically typical condition under the new perceptual circumstance. Only after a significant period for the Earthling to adapt to the new circumstance by accumulating the tokening instances of her orangish experience on Color-inverted Earth can the experience start to appropriately represent things on

Color-inverted Earth as having some properties. A similar situation happens to Color-inverted Earthlings when they come to Earth for the first time. Therefore, Thompson's account predicts a period of perceptual adaptation for one's experience to represent the appropriate property under a novel perceptual circumstance.

However, the prediction seems to result in an awkward consequence. Suppose an Earthling knows that she will inhabit Color-inverted Earth permanently and also clearly knows what will happen to her color experiences on Color-inverted Earth.[6] It seems that she might immediately expect her color experience to be normally caused by properties that are different from what normally causes her color experience on Earth and take the new normal cause as what makes her experiences veridical. It would be odd if she had to live on Color-inverted Earth for a few months (say), waiting for the new properties to become statistically typical so that her experience would start to represent. It seems that typicality is not the only factor for normality. What also matters to the qualification of a normal condition pertains to *what perceptual conditions she takes to be normal* in her reference circumstance, not just what the relevant factors determine.

On the other hand, Millikan (1984) argues that normality derives from the proper function of the biological nature of a reproductively established family. A "reproductively established family" is a technical term that refers to, among other things, a species, a bodily organ type, a syntactic structure type, or a mental state type. Roughly, the proper function of a member of a reproductively established family is the function of the member the successful performance of which is responsible for the continued reproduction of its reproductively established family. For instance, the proper function of a heart is to pump blood in order for it to circulate in the body. Millikan calls an explanation of "how a particular reproductively established family has historically performed a particular proper function" a "Normal explanation" (*Ibid.*: 33).[7] She also uses the term "Normal conditions" to denote the conditions that must hold in order for a proper function to be successfully performed in accordance with a Normal explanation. Thus, a Normal explanation of a heart pumping blood has to tell something about how the heart is made, how it works inside, how the electrical impulses and oxygen supplies are sent to the heart, etc., and a Normal condition is a condition in which these components obtain.

[6] I take the meaning of "knowing" in a very broad sense. In my view, the Earthling need not possess propositional knowledge of Color-inverted Earth. *Understanding* or *grasping* the new environment would suffice.
[7] Millikan capitalizes "Normal" to emphasize that this term should not be read in a statistical sense (*Ibid.*: 34).

Millikan applies this biological picture to accounting for systems of representation by taking them as a sort of reproductively established family. Bee dance that informs nectar and beavers' slap of water that signals danger are two concrete examples. According to Millikan, the proper function of the representing system as a whole is understood in terms of two sub-systems: one *produces* the representation (the producer) and the other *consumes* the representation (the consumer). The function of both sub-systems contributes to the function of the whole system. In particular, dancing (the producer's function) and collecting nectar (the consumer's function) contribute to bees' survival (the function of the whole system); and slapping water (the producer's function) and saving others from danger (the consumer's function) contribute to beavers' survival (the function of the whole system). Furthermore, the consumption system cannot perform successfully unless the surrounding conditions are appropriate. For example, there must be nectar in the location described by the bee dance or a danger in the surroundings. These conditions are the Normal conditions for the proper functioning of the whole representing system, and these Normal conditions comprise the representation's content.

However, Millikan's biological account does not escape the awkward consequence addressed earlier. According to her, the nature of Normal conditions is a historical matter. Since the proper function of a representing system matters to the latter's survival or continued reproduction, it is fixed in terms of how this representing system *qua* reproductively established family has performed successfully in its evolutionary history. As Normal conditions are conditions in which the proper function can be successfully performed, they are also determined by the representing system's evolutionary history. In this way, the content of the Earthling's color experiences is fixed at the Normal conditions on Earth. Although she travels to Color-inverted Earth with her full knowledge of what will be going on with her color experiences, the Normal conditions of her color experience do not henceforth switch to Color-inverted Earth. Thus, if a role representationalist adopts Millikan's account of normality, she must predict that the traveling Earthling suffers from systematic illusion in her color experiences on Color-inverted Earth.

The concern afflicting the naturalist accounts invites a non-naturalist account of normality. I propose that what counts as "normal" is determined by the subject's expectations. Specifically, the "normal" condition in a certain circumstance *C* for a subject *S* just is the condition *S* expects in *C*. Expectations are beliefs about the future. Like beliefs, expectations aim at truth and are evidence-responsive. A belief has a mind-to-world direction of fit, depicting the world as being in

a state of affairs such that the belief is true. If we find new evidence about the world that conflicts with our current beliefs, we will adjust our beliefs to fit the evidence. Similarly, an expectation also has the same direction of fit, depicting the future world as a state of affairs such that the expectation is true. If we find new evidence about the future world, we will adjust our expectations to fit the evidence.

A subject *S*'s expectation fits a certain circumstance *C* only when *S* has sufficient knowledge of *C*. Specifically, for *S*'s expectations to fit the future in *C*, she must know, at least, what *C* is, what typically causes *E* in *C*, and whether she is in *C*. In other words, *S* needs evidence to ground her expectations of *C*. Without knowing these, she might have expectations that are too deviant to fit *C*. If the traveling Earthling does not know Color-inverted Earth, she would falsely expect that her orangish experience will still be caused by C_{600} even if she has been on Color-inverted Earth. This suggests that what counts as "normal" depends on the subject's knowledge of the circumstance.

Taking the non-naturalist account, role representationalism can accommodate the case of interstellar travelers. If an erudite Earthling travels to Color-inverted Earth, C_{500} is the new typical cause of her orangish experience. She will then expect that her orangish experience will be caused by C_{500}. So, C_{500} counts as the normal cause of her orangish experience on Color-inverted Earth. In this way, her orangish experience is leniently veridical when the perceived objects instantiate C_{500} because C_{500} necessitates the instantiation of the role property attributed by her orangish experience's lenient content. By contrast, a naïve Earthling's orange experience is leniently falsidical when she travels to Color-inverted Earth. This is because she does not expect the C_{500} to cause her orangish experience, even if it is the typical cause on Color-inverted Earth. In other words, the object with C_{500} that typically causes the naïve Earthling's orangish experience on Color-inverted Earth does not instantiate the role property of *having a realizer property that normally causes orangish experiences in me on Color-inverted Earth*. Her experience never becomes veridical unless she gets acquainted with the new circumstance, say, by staying on Color-inverted Earth for long as Thompson suggests, or by improving her propositional knowledge of the new circumstance.

5.2.5. Role Representationalism and Two-Tiered Phenomenal Contents

Having explained the key notions of role representationalism, we are in a good position to examine the account as a whole. As I characterized them at the beginning of this chapter, both strict and lenient content are phenomenal content and tightly connected with phenomenology. Although role representationalism is an account of lenient content, this account will be promising if it fits in with a broader picture that accommodates these contents and phenomenology. This sub-section tries to provide an overall picture of the nature and relationships between these contents.

As stated above, the strict content is phenomenal content because it is strictly determined by phenomenology. It represents the properties that the object appears to have, such as phenomenal redness or phenomenal roundness. In addition, I stipulated for the first pass that the lenient content is also phenomenal content because it represents the properties that play the role of producing experiences of a certain sort. Although lenient contents involve components that might not be determined by phenomenology, they can always remain the same if the phenomenology remains the same. I will discuss whether lenient content exactly counts as phenomenal content at the end of this sub-section.

My distinction between strict and lenient content can be accommodated by the two-tiered theory of intentionality. The tenet of the two-tiered theory is that there are two levels of mental content, where the first level is simple representations while the second level is complex representations (Chalmers, 2021: 187ff.; see also Neander, 2017; Mendelovici, 2018, 2019; Williams, 2020; Pautz, 2021). In Mendelovici's and Pautz's accounts, the first-level content is grounded in phenomenology, and the second-level content is *derived* from the associated first-level content. Call the first-level contents the *immediate contents* and the second-level contents the *derived contents*. For example, the phenomenal content <water> is immediate content because it is determined by the watery phenomenal character. Anyone with a watery experience must share <water>. However, different derived contents may be derived from the same immediate content. Oscar and Toscar may derive different derived contents to their complex representations from <water>. Oscar, knowing that H_2O typically causes his watery experience, would derive < H_2O> from <water> through an internal disposition of roughly the form "<water> → < H_2O>"[8]. Similarly, Toscar, knowing that XYZ typically causes her watery experience, would derive <XYZ> from <water> through an internal disposition "<water> → <XYZ>". Thus, the derived contents would differ in Oscar's and Toscar's complex representations.

[8] By "→", I mean the direction of content derivation.

Mendelovici and Pautz have different accounts of the internal disposition that derives derived contents from immediate contents. According to Mendelovici's (2018, 2019) self-ascriptivism, the internal disposition of the form "<*C*> → <*C*+>" (where <*C*> is an immediate content and <*C*+> is a derived content) is characterized by the subject's underlying *cashing out thought* <by <*C*>, I mean <*C*+>> under sufficient reflection. When a subject knows her environment well, then, upon sufficient reflection, she will cash out the immediate content into the derived content that appropriately represents the environment. If the subject is on Earth, she will cash out <water> into <H_2O> upon sufficient reflection through the cashing out thought <by <water>, I mean < H_2O>>. If she is on Twin Earth, she will cash out <water> into <XYZ> upon sufficient reflection through the cashing out thought <by <water>, I mean < XYZ>>.

On the other hand, Pautz (2021) characterizes the underlying disposition by an interpretationist framework:

> Best systems theory: *If*, given one's history of *conscious experiences with immediate content* <C> and consequent *dispositions to act*, all the *best interpretations* assign to one the mental state with derived content <*C*+>, *then* this grounds one having the mental state with <*C*+>. (adapted from Pautz, 2021: 288; italics mine)

The inferential connections between <*C*> and <*C*+> are constrained by whether <*C*+> is one of the best interpretations of <*C*>. Given my experience of a rock flying toward me, I have a reason to believe that a rock is flying toward me. Assuming that a desire to be hit is unreasonable, ascribing such a desire to me renders me irrational. So, based on my experience of the rock and my disposition to shun it, the best interpretation that prevents me from going irrational assigns to me the belief that a rock is flying toward me and the desire not to be hit (Pautz, 2021: 289; see also Lewis, 1986: 38ff., 1994: 427ff.).[9]

The constraint of the inferential connections narrows down the set of possible derived contents from a certain immediate content <*C*>, ruling out all possible contents that bear no inferential connections to <*C*>. Such a constraint prevents content derivation from deriving a

[9] A further question to this approach concerns *where the reasons come from*, as it would be circular if mental contents are grounded in rational inferential roles while the rational inferential roles are also grounded in contents (see Chalmers, 2021: 196ff.). Pautz (2021) nicely addresses this question. This discussion goes beyond the present purpose, so I do not go into the details here.

derived content from an inferentially unconnected immediate content, such as deriving <the view that many VR experiences are leniently veridical> from <water>. Certainly, the constraint does not suffice to single out a definite derived content from the set of possible derived contents. Due to the vagueness of what counts as the best interpretations, it may be indeterminate which definite <*C*+> is derived from <*C*> in a particular situation. For instance, one can derive <C_{700}>, <ripe tomato>, or <apple> from <red>. After all, these derived contents are all contained in the set of the derived contents that are inferentially connected to <red>.

Role representationalism is compatible with both Mendelovici's and Pautz's accounts. I leave open the issue of what is the best account of the internal disposition regarding content derivation. A role representationalist can feel free to adopt either account to supplement my following explanation of the nature and relationship between strict and lenient contents in terms of immediate and derived contents.

Adopting the two-tiered theory, strict contents are subsumed under immediate contents, while lenient contents can be understood as a sort of derived contents. On the one hand, strict contents represent the properties that the object appears to have. Such properties just are what we are immediately aware of when we perceive something, such as phenomenal redness or phenomenal roundness.

Lenient contents, on the other hand, are derived from strict contents. They represent the properties that play the role of producing experiences of a certain sort. According to role representationalism, lenient contents represent the properties that play a specific role, i.e., the role of *having a realizer property that normally causes experiences of a certain sort in the subject under the current perceptual circumstance*. With the clarification of the key notions above, we see that the description of the role is fixed by phenomenology because the description will remain the same if the sort of experiences in question remains the same. Firstly, a role and its realizer are mutually defined, so what particular properties play a role does not matter to its description. Secondly, the "subject" and the "perceptual circumstance" are self-locating components. Their description remains the same in any perceptual experience. Thirdly, the role is individuated in terms of the phenomenal type that its realizer normally causes, so there is no difference in the description of the role without a difference in the experience representing that role. Put them together, the description that singles out specific role properties always stays the same when the sort of experiences remains the same.

According to role representationalism, we derive from the strict content of an experience <*C*> the lenient content <*C+*>, where <*C+*> attributes the role property of *having a realizer property that normally causes experiences of that sort in the subject under the current perceptual circumstance*. In the case of reddish experience, we derive from the strict content <red> the lenient content that attributes the role property of *having a realizer property that normally causes reddish experience in the subject under the current perceptual circumstance*. In the case of roundish experience, we derive from the strict content <round> the lenient content that attributes the role property of *having a realizer property that normally causes roundish experience in the subject under the current perceptual circumstance*.

Strictly speaking, lenient content is not a sort of phenomenal content, which is defined as the content determined by phenomenology. Since lenient content is derived from strict content, and its derivation is not fully constrained by its inferential connections to strict content, it is not built into the subject's perceptual phenomenology. Intuitively, it is possible for a subject to take her experience of a certain sort to mean something that is different from the role properties identified by role representationalism. In particular, it is possible for a subject to cash out from the set of the derived contents that are inferentially connected to <red> a lenient content that is distinct from the one that attributes the role property of *having a realizer property that normally causes reddish experience in the subject under the current perceptual circumstance*. In this sense, lenient contents are not completely determined by phenomenology, so they do not count as phenomenal contents.

Although lenient contents are not properly phenomenal contents, they are nonetheless interestingly similar to phenomenal contents. According to role representationalism, lenient contents are determined by phenomenology *plus* an internal disposition. Indeed, the internal disposition is constrained by the subject's knowledge about her environment and, in turn, by the external world. However, the internal disposition that partly determines lenient contents remains the same across different kinds of perceptual experiences. Role representationalism claims that lenient contents attribute the role properties whose description is ubiquitous in experiences of all kinds, including color, size, shape experiences, and even experiences in other modalities. That is, for any kind of experience, lenient contents always attribute the role property of *having a realizer property that normally causes experience of that sort in the subject under the current perceptual circumstance*. So, we can say that the lenient contents of an experience always follow its

phenomenology. Although lenient contents are not properly phenomenal, they are *weakly* phenomenal.

To sum up, this section laid out role representationalism and articulated its key notions. I also pointed out the advantage of this account when combined with the two-tiered theory of intentionality. The next section will examine how this broad picture coheres with virtual veridicalism.

5.3. Role Representationalism and Virtual Veridicalism

The preceding section proposed role representationalism as a detailed picture of the contents of perceptual experience. This section further assesses the coherency of this account with virtual veridicalism. Specifically, I examine whether role representationalism can accommodate all the constraints of virtual veridicalism listed in Section 5.1. Virtual veridicalism is a view pertaining to lenient veridicality, and role representationalism is an account of lenient content. If my argument succeeds, role representationalism and virtual veridicalism are mutually supported.

Recall *Virtual Basketball* discussed in Section 3.4. Suppose that when Sally sees a physical basketball in an ordinary environment, she sees the physical basketball as orange, α-sized, and round. When she sees a virtual basketball in a perfect virtual world (i.e., a virtual world that *completely* screens off non-virtual things and then appears the same as non-VR environments), she also sees the virtual basketball as orange, α-sized, and round. Call a token of the group of her experiences while seeing the physical basketball E_P and a token of the group of her experiences while seeing the virtual basketball E_V. Sally cannot distinguish between E_P and E_V through introspection. This section assesses the consistency of role representationalism and virtual veridicalism by examining whether it can accommodate the case of E_P and E_V in a way that does not violate any virtual veridicalist constraint. This case serves as an instance of the general case involving a pair of subjectively indistinguishable VR and non-VR experiences.

First, *Same Phenomenology* follows from subjective indistinguishability between E_P and E_V. Representationalists typically assume the common kind assumption, i.e., subjective indistinguishability entails phenomenal identity. So, we say that E_P and E_V share the orangish, the α-size-ish, and the roundish phenomenal character. According to role representationalism, they also share the same lenient content. For ease of discussion, I focus on the orangish character and its lenient content.

Second, given *Veridical Non-VR Experience*, the role property that E_P's lenient content attributes should be instantiated in ordinary environments so that E_P can be veridical. Likewise, *Veridical VR Experience* requires that the role property that E_V's lenient content attributes be instantiated in virtual worlds. Thus, the property that both E_P's and E_V's lenient contents attribute must be instantiated in both virtual and ordinary environments.

Third, *Different Reliable Cause* requires that the properties that reliably cause the group of experiences that contains E_P and the group of experiences that contains E_V are categorically different. In ordinary environments, the group of E_P is reliably caused by C_{600}. In virtual worlds, the group of E_V is reliably caused by a virtual property distinct from C_{600} (whatever it is, call it VC_{600}).

Put them together, these constraints require role representationalism to account for the theses that 1) E_P's and E_V's lenient contents attribute the same property that is instantiated in these environments and 2) E_P and E_V track distinct properties in their respective environments.

Role representationalism contains sufficient resources to accommodate the theses. Accordingly, E_P's and E_V's lenient contents attribute the role property of *having a realizer property that normally causes orangish experience in me under the current perceptual circumstance*. Since C_{600} normally causes Sally's orangish experience in ordinary environments, and VC_{600} normally causes Sally's orangish experience in virtual worlds, the role property that E_P's and E_V's lenient contents attribute is realized by C_{600} in the ordinary environment and by VC_{600} in the virtual world. *Same Phenomenology* and *Different Reliable Cause* are accounted for in this way. Moreover, since C_{600} is instantiated by the physical basketball in the ordinary environment, and VC_{600} is instantiated by the virtual basketball in the virtual world, the role property is also instantiated under both circumstances. *Veridical VR Experience* and *Veridical Non-VR Experience* also hold.

Fourth, role representationalism satisfies these constraints not only in perfect virtual worlds but also in mixed worlds (i.e., virtual worlds that *partly* screen off non-virtual things and then appear the same as non-VR environments). Now, suppose that E_V occurs in a mixed world that has a 50%-50% reality-virtuality (RV) mixture. The change of E_V's circumstance does not affect the satisfaction of *Same Phenomenology*, *Different Reliable Cause*, and *Veridical Non-VR Experience*. As E_V does not change, it still has the same phenomenology as E_P does. E_V is still caused by VC_{600}, which is categorically different from what causes E_P, *viz*. C_{600}. The change does not affect the veridicality of E_P, either. Moreover, role representationalism predicts that E_V is leniently veridical

in the mixed world. E_V's lenient content attributes the role property of *having a realizer property that normally causes orangish experience in me under the current perceptual circumstance*. As VC_{600} is in fact a typical cause of Sally's orangish experience in the mixed world, it is natural for Sally, with sufficient knowledge, to expect VC_{600} to cause her orangish experience in the mixed world. Since the instantiation of VC_{600} necessitates the instantiation of the role property attributed by E_V's lenient content, E_V is leniently veridical in the mixed world in that its object (*viz.*, the virtual basketball) instantiates the property that E_V's lenient content attributes. *Scope* is satisfied in a mixed world with a 50%-50% RV mixture.

To generalize, *Scope* will be satisfied in a mixed world with different percentages of RV mixture if the cause of Sally's orangish experience counts as its normal cause. Recall that a normal cause of a certain experience in a subject under a circumstance is what is expected to cause that experience in that subject under that circumstance. Provided that Sally has sufficient knowledge of any mixed world, she will expect C_{600} or VC_{600} to cause her orangish experience. Now, the question is, depending on the ratio of RV mixture, to what extent is it reasonable for Sally to expect C_{600} or VC_{600} to cause her orangish experience? At first glance, it seems unreasonable to still expect that VC_{600} will cause her orangish experience in a mixed world with an 80%-20% or a 90%-10% RV mixture.

My answer: It depends on the probability value the subject assigns to the given mixed world. It seems intuitive that if Sally's orangish experience occurs in a mixed world with a 60%-40% RV mixture, she will reasonably expect VC_{600} as well as C_{600} to be its cause. It might still be the case when the mixed world contains a 65%-35% or a 70%-30% RV mixture. But Sally might hesitate to expect VC_{600} to cause her orangish experience in a mixed world with an 80%-20% RV mixture because the probability is too low. Of course, the extent to which an expectation is reasonable is vague, and I do not think there is a definite solution. At least, we can safely say that VC_{600} reasonably counts as a normal cause of Sally's orangish experience in any mixed world where the ratio of RV mixture is distant from 1. *Scope* is still satisfied in these mixed worlds.

As an extra advantage, role representationalism allows an experience to have more than one normal cause under a circumstance. For instance, in a mixed world where the RV mixture ratio is distant from 1, the subject can reasonably expect her experience to have plural normal causes. In particular, it seems intuitive to say that if Sally's orangish experience occurs in a mixed world, say, with a 50%-50% or a 60%-40% RV mixture, she will reasonably expect her experience to be

caused by C_{600} or VC_{600} or both. If so, both C_{600} and VC_{600} realize the role property attributed by the lenient content of her orangish experience. In this way, Sally's orangish experience can be leniently veridical in the mixed world.

I end this section by addressing an implication that role representationalism provides for our perception in daily usage of VR. Although a VR user may be extremely familiar with virtual worlds, she has to leave VR and frequently switch between virtual and non-virtual environments in practice. Our intuition would be that once the user is familiar with a virtual world, her experiences of objects in the world are supposed to be veridical even though she left and re-entered it. Role representationalism can maintain stable veridicality across virtual and non-virtual environments. Provided the user's knowledge of the virtual world, the user will expect a cause in a virtual world that is different from the cause in the ordinary environment. In this way, even if Sally frequently switches between virtual and non-virtual circumstances, her experience does not take time to (re-)adapt to fix what its lenient content attributes under each circumstance. I believe virtual veridicalists would be happy to accept this implication. As VR users indeed frequently switch between virtual worlds and ordinary environments, it would be theoretically fruitful to preserve perceptual veridicality independently of the length of perceptual history in each circumstance.

In conclusion, role representationalism coheres with virtual veridicalism. Specifically, the view satisfies all the constraints of virtual veridicalism, i.e., *Veridical VR Experience*, *Veridical Non-VR Experience*, *Same Phenomenology*, *Different Reliable Cause*, and *Scope.* Therefore, role representationalism and virtual veridicalism are mutually supporting positions. It is plausible that both are simultaneously true.

5.4. Objections and Responses

So far, I have put forward role representationalism and argued that it coheres with virtual veridicalism. Some might be concerned about whether role representationalism itself faces problems that are more easily solved by its alternatives. Presumably, if role representationalism cannot resolve these problems, it is doubtful to take the account as a sufficiently good account of perceptual experience. This section defends role representationalism from a few concerns.

5.4.1. Why Not Realizer Representationalism?

Given role representationalism, one might wonder if a similar view, realizer representationalism, would be a better alternative. This sub-section compares their difference and rejects the realizer alternative.

Like role representationalism, realizer representationalism consists of two parts. It agrees with role representationalism on *Representationalism*. But it disagrees with role representationalism on content attribution. Realizer representationalism accepts *realizer functionalism about lenient content*:

> *Realizer functionalism about Lenient content*: The lenient content of a perceptual experience with phenomenal character *C* is a proposition of the form <something is *F*>, where *F* is the property that plays a role of *normally causing experiences with C in the subject under the current perceptual circumstance*.

Similarly, realizer functionalism can construe the key notions — *subject*, *perceptual circumstance*, and *normal condition* — in the same way as role functionalism does. In this respect, role functionalism shares with realizer functionalism all the theoretical advantages I argued for in Section 5.2.

The difference between role functionalism and realizer functionalism rests on the difference in what *F* attributes. Role functionalism takes *F* to attribute the higher-order role properties, while realizer functionalism takes *F* to attribute the first-order realizer properties. According to role functionalism, *F* is the functional role that is constituted in terms of its relation to subjects. By contrast, realizer functionalism takes *F* as something that *happens* to play such a role, so the functional role is not a constituent of *F*. Instead, it is what singles out *F* in a certain perceptual circumstance. For instance, C_{600} plays the orangish experience role on Earth, so realizer functionalism predicts that orangish experience represents C_{600}. But C_{600} does not constitute the orangish experience role. It just *happens* to play the role in the very circumstance. If something other than C_{600} happens to play the role in the same circumstance, realizer functionalism will predict that orangish experience does not represent C_{600} but the new property.

Given the difference, one might wonder why I defend role representationalism instead of realizer representationalism. This sub-section offers two reasons for my preference.

My reasons for role representationalism follow Cohen's (2009: §7) arguments for role functionalism of color. The first reason concerns the explanatory demand of *commonality*. Consider Yablo's discussion about property attribution:

> "Here is the scenario. There is a predicate '*F*' and we have various things we want to say with it — things we regard as quite like true, or even certainly true. But we are troubled. Granted that '*F*' applies to roughly the objects we suppose, what property or relation or condition or what have you does it apply *in virtue of*? Information about the property is not lacking; there are the various statements we are inclined to make using '*F*', and these, being presumably correct, add up to a considerable data base. What bothers us is that the information feels *circumstantial*; we learn what the property is *like*, not what it *is*. The urge thus arises to *identify F*-ness, to single it out and elucidate its nature. Only then will we know what makes our predications true." (Yablo, 1995: 477–8; original emphasis)

To explain what grounds our accurate attribution of *F* to various objects, the best answer is that these objects share *F*-ness. That is, the attribution of *F* to these objects is true because they share *F*-ness. However, this explanation only tells what *F*-ness *does* (e.g., which objects have it and which do not). It does not tell what *F*-ness *is*. To answer the latter question, we need to explain why *F*-ness is shared by all and only *F*-objects (see also Cohen, 2009: 188). For example, we attribute "orange" to orange objects. If such an attribution is true, these objects do share *orange-ness*. However, the fact that these objects share orange-ness does not tell anything about the nature of orange-ness. It only tells which objects have such a common nature. To elucidate the nature of orange-ness, it is required to explain why orange-ness is shared by all and only orange objects.

This commonality constraint demands role functionalism to explain the commonality of all and only *F*-objects. Take role functionalism of "can-opener" as an example. Suppose that "can-opener" is functionally defined. A can-opener may be made of wood, plastic, metal, etc. To explain the commonality between these physically heterogeneous instances, role functionalism takes the can-opener role as the common nature shared by these physical properties and attributes a role property to that nature, *viz., being a can-opener*. In doing so, role functionalism explains the nature of can-openers by virtue of the role property *being a can-opener* that they commonly share. Every

object that has the property is a can-opener, and any object that does not have such a property is not a can-opener.

Similarly, role functionalism about lenient content takes *F* to attribute role property to the perceived objects. Since the attributed property is the functional role that is constituted in terms of its relation to the perceivers, it is necessary that the property produces the relevant effect on the perceivers while being perceived. It is guaranteed that every property attributed by the lenient content of orangish experience plays the orangish experience role, and any property that is not attributed by the lenient content of orangish experience does not play the orangish experience role.[10] Role functionalism explains the common nature of orange objects by virtue of their having the very role property. Thus, the commonality constraint is met by role functionalism.

However, realizer functionalism does not meet the commonality constraint. The view takes *F* to attribute the properties that *happen* to play the relevant role to the perceived objects, so it hardly explains why all and only these properties share *F*-ness. Recall the "can-opener". Realizer functionalism tries to reduce the property *being a can-opener* to the properties that play the can-opener role, namely *being wood*, *being plastic*, *being metal*, etc. This view does not take these properties to be constituted in terms of the can-opener role: They could have not played the role. But the commonality constraint also demands the view to explain why all and only these properties play the can-opener role. Since these properties do not suffice for a can-opener, it is hard for realizer functionalism to find a plausible unifying explanation of why all and only the physically heterogeneous properties count as can-openers.

The same problem goes to realizer functionalism about lenient content, as it takes *F* to attribute the properties that *happen* to play the relevant role to the perceived objects. The lenient content of orangish experience attributes the properties that happen to play the orangish experience role. Because it is possible for these properties not to play the same role, the nature of the functional role cannot be explained by the objects' having of these realizer properties. Realizer functionalism does not guarantee that all and only properties that the lenient content of orangish experience attributes play the orangish experience role. So, this account does not meet the commonality constraint.

[10] Notice that the orangish experience role involves the constraint of normality. Not all instances in which orangish experience is instantiated occur in normal condition. So, the statement is compatible to the possibility of misperception.

In the literature, realizer functionalists offer several ways to meet the demand for commonality, but they all fail. First, they might explain it by pointing out that the heterogeneous realizer properties belong to a common disjunction (Jackson, 1998; McLaughlin, 2003). But this explanation is unsatisfactory, for it does not explain why these properties *necessitate* a role in common, but only the contingent fact that these properties *happen* to play the role. Second, realizer functionalists might accommodate the commonality by the common causal analysis of the functional roles. Since a functional role is analyzed in terms of its bearing certain causal relations to its causes, effects, and other roles, the fact that a set of distinct properties plays a certain role implies that they bear these causal relations *in common*. However, such an answer dissatisfies the explanatory demand, since it merely changes the topic from asking why all and only the diverse realizers play the common role to asking why they share these causal relations. As Cohen (2009: 189) diagnoses, this solution fails unless the role itself is taken as constitutive of the shared commonality. But this move shifts to role functionalism. The third alleged solution is to ground the commonality in a metaphysical thesis, such as "Necessarily all and only the orange things play the orange role." For this thesis to be informative, the "orange things" should be read *de re*; otherwise, the *de dicto* reading is trivially true by the definition of realizer and role. Realizer functionalism is nonetheless unable to warrant the truth of the thesis in *de re* reading, as it must make room for the possibility in which an orange realizer (*de re*) does not play the orange role. Therefore, these responses do not prevent realizer functionalism from the difficulty in explaining the demand for commonality.

My second reason against realizer functionalism concerns the *metaphysical necessity* of the properties attributed by lenient contents. Content attribution affects how we refer to external properties through linguistic expressions. Presumably, the properties that color content attributes are colors, and the properties that shape content attributes are shapes, etc. So, if role representationalism is true, colors and shapes *are* the role properties. In contrast, if realizer representationalism is true, colors and shapes *are* the realizer properties. The difference between the two views results in different predictions of the metaphysical necessity of colors and shapes.

Presumably, all phenomenal characters are necessarily different from others. An orangish character is necessarily different from a reddish character because they are distinct what-it-is-likeness. An orangish character cannot be felt like a reddish character because every character that is felt like a reddish character *is* a reddish character, not anything else. Similarly, a roundish

character is necessarily different from an ellipse character. An account of perceptual experience is more desirable if it predicts the necessary difference in the properties represented by experiences with necessarily distinct characters.

The prediction made by role representationalism fits the metaphysical necessity in phenomenology. Accordingly, the represented property is the functional role that is constituted in terms of its relation to experience, so it is singled out by the experience type in relation to it. For instance, the lenient content of orangish experience attributes the property regarding the orangish experience role, while that of reddish experience attributes the property pertaining to the reddish experience role. As their role is singled out by the experience type, and since orangish experience is necessarily different from reddish experience, the orangish experience role is also necessarily different from the reddish experience role. In turn, the property defined in terms of the orangish experience role (i.e., what we call "orange") is necessarily different from the property defined in terms of the reddish experience role (i.e., what we call "red"). Analogously, the property defined in terms of the roundish experience role (i.e., what we call "round") is necessarily different from the property defined in terms of the ellipse-experience role (i.e., what we call "ellipse"). Therefore, role representationalism predicts a metaphysically necessary difference in the external properties attributed by the lenient contents of necessarily distinct kinds of experience.

However, realizer representationalism does not predict such a metaphysical necessity in phenomenology. As the view takes the property attributed by lenient contents as what *happens* to play the functional role of producing an experience type, the property is not constituted by the role it plays. In particular, the lenient content of orangish experience attributes the property that plays the orangish experience role, so what we call "orange" is the properties that contingently play the role. In this respect, it is possible for a represented property not to play the role it actually plays. In other words, it is possible for the "orange" property not to play the orangish experience role. Suppose the "orange" property is C_{600} on Earth. The way that realizer representationalism picks out what lenient contents attribute does not prevent C_{600} from playing another role. If what we call "red", C_{700}, satisfies the orangish experience role, then realizer representationalism predicts the property to be attributed by the lenient content of orangish experience. Hence, realizer representationalism must deny the claim that the orange property is necessarily different from the red property. The denial of the necessity conflicts with the necessary difference between orangish

and reddish experiences, suggesting that realizer representationalism makes a wrong prediction about what lenient contents attribute.

Given the reasons above, role representationalism is preferable to realizer representationalism. Role representationalism is resourceful to meet the demand of commonality by attributing high-order properties to objects that play a certain common role. The view also rightly predicts the necessary difference between the represented properties. Realizer representationalism, by contrast, fails to do both. So, role representationalism is a better view of perceptual experience than realizer representationalism.

5.4.2. Are Role Properties Causally Efficacious?

This sub-section turns to a challenge to broadly role functionalism. In the literature, role functionalism of mind is criticized for implying a metaphysically necessary consequence that no role property is causally efficacious (Rupert, 2006; see also Robb et al., 2023: §6.6; Bennett, 2007: §7). As role functionalism of mind is motivated to preserve mental causation, the consequence is fatal to this view. Analogously, since role representationalism commits to role functionalism about lenient content, according to which lenient contents attribute functionally defined role properties, this view might potentially result in a similar consequence that these role properties are causally inefficacious. Let me elucidate.

Rupert (2006) argues against the causal efficacy of role properties in terms of the *problem of metaphysically necessary effects*. According to role functionalism of mind, a mental state is a state that is completely defined in terms of its causal relations involving inputs, outputs, and other mental states. If so, when a causal claim about a mental state is made, it is vacuously true. For example, if pain is functionally defined as a higher-level state of being in the lower-level state that causes wincing in certain circumstances, then the causal claim in the form that "pain causes wincing" is uninformative, since the state in question would not be pain if it did not cause wincing. Generalizing the case, every causal claim about a functionally defined mental state can be translated into the form that "the state that causes *e* causes *e*". Since causing *e* is a constituent of the state's functional role, such a causal claim is metaphysically necessary and hence vacuously true. The consequence casts doubt on the causal efficacy of functionally defined mental states. If mental states can be causes, they metaphysically necessitate their effects. But because causation

cannot be metaphysically necessary, role functionalists are forced to conclude that mental states are causally inefficacious and deny mental causation.

Some might apply Rupert's argument to reject role representationalism, arguing that the view results in a similar problem of metaphysically necessary effects of the role properties that lenient contents attribute. These properties are completely defined in terms of their causal relations to a certain experience type. In particular, the lenient content of orangish experience attributes the role property of *having a realizer that normally causes orangish experiences in the subject under the current perceptual circumstance*. So, when we make a causal claim, say, "an orange ball causes orangish experience", such a claim is metaphysically necessary and hence uninformative. This is because the underlying form of the causal claim is something like "a ball having the role property of *having a realizer that normally causes orangish experiences in the subject under the current perceptual circumstance* causes orangish experience". As a result, role representationalists also have to deny the causal efficacy of the role properties.

I think this objection is right, since the problem of metaphysically necessary effects cannot be avoided by role representationalism. But I do not think the problem is fatal to the account, as it does not refute role representationalism but only challenges the assumption that the role properties are more or less causally efficacious. In fact, role representationalism need not take the role properties to have any causal power. As discussed above, it is a common mistake to identify what an experience represents with what it tracks, for the identification wrongly predicts unwanted materials in the representation but also wrongly omits essential materials from the representation (Mendelovici, 2018: Ch. 3). Also, the motivation to preserve the causal efficacy of the represented properties rests on our aversion to epiphenomenalism (e.g., Kim, 1998). But as Segal (2009) argues, there is no problem with the epiphenomenalist consequence that functional roles are not causally efficacious, since they still "play indispensable roles in excellent causal explanations, just like their physical counterpart [like stiffness and impact strength]" (*Ibid.*: 101). A stone's stiffness and impact strength are functionally defined role properties that are realized by its material constitution, so they are causally inefficacious. But they are still causally explanatory with respect to the stone breaking the window in that they feature in the relevant explanation. The explanation that the stone breaks the window because it is stiff and strong is intuitively correct. It would be insufficient to explain the event solely by the stone having a certain material constitution. Analogously, a ball being "orange" is causally explanatory with respect to our having orangish experience while seeing

the ball. Intuitively, we explain the having of orangish experience in terms of the ball being "orange". It would be insufficient to explain it by the ball having C_{600}. Hence, even if role properties are not causally efficacious, we cannot avoid citing "the role properties in causal explanation of the effects that comprise their roles so long as those role properties are causally relevant to those effects" (Cohen, 2005: 132).

In sum, the problem of metaphysically necessary effects does not threaten role representationalism, since this account does not take the role properties as causal properties. Instead, the properties are instantiated by virtue of the instantiation of their realizers that do the causal jobs. So, it is theoretically adequate that the role properties are not causally efficacious but only play a crucial role in the relevant causal explanation.

5.5. Conclusion

In conclusion, I proposed in this chapter role representationalism as a new account of perceptual experience and the lenient content, arguing that it coheres well with virtual veridicalism. I also defended role representationalism from a few objections by arguing that the view can resolve the problems faced by similar accounts. Therefore, role representationalism and virtual veridicalism are mutually supporting: Virtual veridicalism offers a desideratum for an account of perceptual experience that is sufficiently flexible to accommodate perceptual veridicality in both ordinary and virtual environments, and role representationalism provides a coherent picture of lenient content to underpin the desideratum. Together with the conclusion of Chapter 4, I conclude that role representationalism is the best account that provides a coherent picture of content in support of virtual veridicalism.

Chapter 6: Chalmers on Virtual Veridicalism

So far, I have completed my defense of virtual veridicalism. To briefly review, Chapter 2 argued that virtual objects exist and instantiate virtual properties, suggesting that something is perceptible in VR. Chapter 3 argued for a specific version of virtual veridicalism, contending that many VR experiences are leniently veridical. Chapters 4 and 5 looked for a reasonably detailed picture of the contents of perceptual experiences that coheres with this version. While Chapter 4 contented that no currently prominent representationalist accounts are consistent with this version, Chapter 5 proposed a new account — role representationalism — and argued that it perfectly coheres with it. Thus, the version of virtual veridicalism defended in Chapter 3 and role representationalism are mutually supporting.

This chapter turns to assessing other defenses of virtual veridicalism. I target Chalmers's (2017) argument, which is the sole defense of this view in the literature so far. This chapter argues that his argument does not succeed in defending virtual veridicalism unless it presupposes role representationalism as an underlying theory of content. Section 6.1 presents and analyzes Chalmers's argument. Section 6.2 disambiguates its key notion — cognitive orientation — and argues that it should be interpreted as a specific sort of cognitive penetration, i.e., diachronic cognitive penetration. Section 6.3 argues that although Chalmers's argument becomes reasonably clear under this interpretation, it is inconsistent with virtual veridicalism. I also diagnose that the inconsistency can be resolved if an appropriate account of perceptual content underpins Chalmers's defense. Section 6.4 offers my suggestion for resolution: Chalmers's defense better presupposes my account — role representationalism — as the underlying theory of content.

6.1. Chalmers's Defense of Virtual Veridicalism

As addressed in Section 2.1, Chalmers (2017) defends *weak virtual digitalism*, the view that virtual objects exist and depend upon distinct, existing data structures. He argues that virtual objects instantiate two types of properties: A property *X* is a *virtual-inclusive property* when *virtual X* is identical to *X* (e.g., a virtual library is a library), while it is a *virtual-exclusive property* when *virtual X* is not identical to *X* (e.g., a virtual kitten is not a kitten; see Section 2.3). Chalmers contends that some low-level perceptible virtual properties (i.e., virtual colors, shapes, and sizes) are virtual-exclusive properties. For instance, virtual redness is not the property to which we

ordinarily ascribe reddish phenomenal character, i.e., physical redness. Rather, virtual redness is considered "either as the power to cause reddish experiences in normal VR conditions, or as the property that normally causes reddish experiences in those conditions" (*Ibid.*: 322). Either way, virtual redness can be understood as what plays the role of normally causing reddish experiences in a virtual world. Similarly, we understand virtual square-ness as what plays the role of normally causing squarish experiences in a virtual world, so do we understand virtual six-foot-tall-ness. In short, virtual colors, shapes, and sizes are not identical to physical colors, shapes, and sizes.

Chalmers's defense of virtual veridicalism is also committed to all the constraints I laid out in Section 4.1. Following my terminology in previous chapters, I mean this specific version by "virtual veridicalism".

1. *Veridical VR Experience.* Many VR experiences are leniently veridical.
2. *Veridical Non-VR Experience.* Many non-VR experiences are leniently veridical.
3. *Same Phenomenology*. A VR experience and its corresponding non-VR experience sometimes share some phenomenal characters.[1]
4. *Different Reliable Cause.* The properties that reliably cause a VR experience are different in kind from those that reliably cause its corresponding non-VR experience.[2]
5. *Scope. Veridical VR Experience*, *Veridical Non-VR Experience*, *Same Phenomenology*, and *Different Reliable Cause* are satisfied in both perfect virtual worlds and mixed worlds.[3]

In the target paper, Chalmers defends the view that many VR experiences are as veridical as ordinary experiences, which commits him to *Veridical VR Experience* and *Veridical Non-VR Experience*. Also, he allows different properties that play the same experiential role in different perceptual circumstances to be represented by the same type of experience. For instance, reddish experiences, while representing physical redness in ordinary environments, represent virtual redness in virtual worlds, as "an object is *virtually red* when it produces reddish experiences in the

[1] To repeat, by VR experience, I mean a *group* of token experiences of a certain type that occur in a virtual world; and by non-VR experience, I mean a *group* of token experience of a certain type that occurs in an ordinary environment.
[2] A VR experience is the corresponding experience of a phenomenally identical non-VR experience and vice versa.
[3] A perfect virtual world is a virtual world where all entities that appear to the subjects are virtual entities, while a mixed world is a virtual world where some entities that appear to the subjects are virtual entities and the others are physical entities.

conditions that are normal for virtual reality" (*Ibid.*: 322). In brief, physical redness and virtual redness normally produce reddish experiences in their respective environments, and they are distinct properties in kind. Hence, Chalmers is also committed to *Same Phenomenology* and *Different Reliable Cause*. At last, Chalmers contends that perception of virtual objects in mixed worlds need not involve illusions since "virtual objects in [mixed worlds] have the same ontological status as those in [perfect virtual worlds]: they are digital objects and are perfectly real" (*Ibid.*: 344–5). In a nutshell, he thinks that there is no fundamental difference in our perception of virtual objects in perfect virtual worlds and in mixed worlds, and the aforementioned constraints still hold in mixed worlds (*viz. Scope*).

In addition, Chalmers extensively defends Fregean representationalism (the Chalmers-Thompson account, as I labeled in Section 4.4) in other papers (Chalmers, 2004, 2006).[4] Although he is silent about this view in the target paper, I assume that he still endorses Fregean representationalism in his defense of virtual veridicalism. So, this chapter still operates within the representationalist framework.

Chalmers's commitment to the five constraints above and Fregean representationalism is central to my assessment of his defense of virtual veridicalism. I will examine the consistency of his defense with his commitment to them.

Chalmers defends virtual veridicalism on the grounds that some VR users can veridically perceive virtual properties due to the process of *cognitive orientation*. In his coinage of the notion, "cognitive orientation" refers to the influence of cognition on perception that occurs in *sophisticated* perceivers, as opposed to *naïve* perceivers. A sophisticated perceiver is familiar with the nature of the things she perceives, whereas a naïve perceiver has no idea about their nature. According to Chalmers, virtual veridicalism is true in that sophisticated VR users perceive virtual objects as having the virtual properties that those objects instantiate. They can do so because their background knowledge that they are in VR *orients* the perceptual content from representing physical properties to representing virtual properties. By contrast, naïve VR users, who have no knowledge that could orient the perceptual content, always misperceive virtual objects. They always perceive virtual objects as having physical properties these objects do not have. Here is Chalmers's argument:

[4] For detailed exposition of this view, please see Section 4.4.1.

P1: Sophisticated VR users, when using VR, know that the colors, locations, and shapes (CLSs, for short) they perceive are virtual CLSs.

P2: If P1, then that knowledge orients the contents of the users' VR experiences, such that they perceive virtual objects as having virtual CLSs.

C1: Therefore, sophisticated VR users perceive virtual objects as having virtual CLSs.

P3: Virtual objects do have virtual CLSs.

C2: Therefore, many VR experiences are veridical. (*Ibid.*: §5).

P1 is an empirical fact that those familiar with VR devices know they are seeing virtual things while using the devices. Naïve users, in contrast, have no such knowledge. I believe there is no reason to deny P1. P3 is Chalmers's metaphysical view of virtual objects and properties. According to Chapter 2, P3 is plausibly true. As virtual objects reliably cause color, space, and shape experiences in VR users, it is reasonable to say they have virtual CLSs. The consequent of P2 is cognitive orientation in sophisticated VR users' perception. Since it is unclear how P2 is the case, the premise needs more clarification and support. Chalmers argues for P2 by analogy:

P2.1: If sophisticated drivers, when driving, know that the things seen in the rear mirror are located *behind* their car, then that knowledge orients the content of their experiences, such that they perceive the things in the rear mirror as being located *behind* the car.

P2.2: If P2.1, then P2.

Therefore, P2.

P2.1 is cognitive orientation in sophisticated drivers' perception. P2.1 is similar to P2, but the case addressed by the former is much more common than the latter. For ease of discussion, I take P2.1 for granted.

P2.2 is an analogical inference drawing from the driver case to the VR user case. To justify the analogy, Chalmers argues that sophisticated drivers typically perceive things in the rear mirror *as located behind* due to their background knowledge that the mirror is present. Naïve drivers have no such knowledge, so they typically perceive things *as located in the front*. He argues that perceptual content differs between sophisticated drivers and naïve drivers because "background knowledge [of sophisticated drivers] helps *orient* one to the perceived world, giving a global

interpretation to what is perceived" (*Ibid.*: 330). According to Chalmers, the occurrence of cognitive orientation rests on the satisfaction of four key features:

1) *Knowledge*: Sophisticated drivers know the mirror is present.
2) *Familiarity*: Sophisticated drivers are familiar with how to use mirrors and how mirrors function in perception.
3) *Action-dependence*: Some pattern of sophisticated drivers' actions depends on the interpretation of mirror perception.
4) *Naturalness*: The interpretation on which the things are located behind is more natural than the interpretation on which they are located in the front.

Although these features are not necessary conditions for cognitive orientation, it is reasonable to say that the driver case involves cognitive orientation if more features are satisfied.

Chalmers argues that the way the VR users are affected by these features is the same as that in the driver case. So, given that cognitive orientation occurs in the driver case, it also occurs in the VR user case. To justify this claim, he contends that the VR user case also satisfies the four features.

1) *Knowledge*: Sophisticated VR users know they are in VR.
2) *Familiarity*: Sophisticated VR users are familiar with how VR functions in perception.
3) *Action-dependence*: Some pattern of sophisticated VR users' actions depends on the interpretation of VR perception.
4) *Naturalness*: The interpretation on which things are virtual is more natural than the interpretation on which they are physical.

Provided that these features are satisfied in the VR user case, Chalmers concludes that P2.2 is true. Naïve VR users perceive virtual objects as physical. But after they learn that they are using VR, they will "act in ways that turn on interpreting themselves to be in VR" (*Ibid.*: 331). Eventually, due to cognitive orientation, they become sophisticated and perceive virtual objects as virtual.

Notably, Chalmers explicitly addresses the perceptual experiences of CLSs in the target paper. I assume that the targeted sort of perceptual content is *phenomenal content*, and that his

argument targets cognitive orientation specifically in the phenomenal content of CLSs. In brief, I assume that Chalmers endorses the phenomenal content view (i.e., for any experience with low-level phenomenal characters, there are contents determined by these characters).

6.2. Cognitive Orientation and Cognitive Penetration

As we have seen, the critical premise of Chalmers's argument is P2.2. To assess P2.2, we must know whether the analogical inference about cognitive orientation is reasonable. Unfortunately, the meaning of cognitive orientation remains unclear in the target paper, hindering the overall assessment of Chalmers's argument. Hence, this section aims to clarify what cognitive orientation should be.

6.2.1. Synchronic vs. Diachronic Cognitive Penetration

Chalmers argues for the occurrence of cognitive orientation in the VR user case in terms of its occurrence in the driver case. He defines cognitive orientation as a sort of *cognitive penetration*, by which he means that "cognition influences perception" (*Ibid.*: 330). However, it is highly controversial about what counts as a cognitive influence on perception in the relevant literature (Zeimbekis & Raftopoulos, 2015). There are many ways that cognition can influence perception. Different conceptions of cognitive influence result in different interpretations of cognitive penetration. Thus, it should also be clarified which interpretation of cognitive penetration Chalmers means. This sub-section presents two interpretations of the notion. The following sub-sections will examine which one fits Chalmers's notion of cognitive orientation.

Cognitive penetration is a more qualified notion than cognitive orientation. In the literature, many cases where cognition influences perception do not count as cognitive penetration. For instance, Siegel proposes a criterion of cognitive penetration that excludes some sorts of cognitive influence:

> "If visual experience is cognitively penetrable, then it is nomologically possible for two subjects (or for one subject in different counterfactual circumstances, or at different times) to have visual experiences with different contents while seeing and attending to the same distal stimuli under the same external conditions, as a result of differences in other cognitive (including affective) states." (Siegel, 2012: 205–6)

Cognitive penetration occurs only when cognition influences perceptual content despite most identifiable perceptual factors (e.g., attention, distal stimuli, external conditions, etc.) being fixed. A perceptual experience is not penetrated if its content changes due to any difference in these factors. Let me explain Siegel's criterion with three examples.

> **1a. Coffee.** You are reading a book. When you desire coffee, the belief that coffee is available in the kitchen motivates you to put the book on the desk and go to the kitchen. Because of the influence of your desire and belief, the content of your visual experience changes from the book to the scenery along the way to the kitchen.

Apparently, this is a case where cognition influences perception. In **1a**, perception is influenced by cognition in the way that your belief influences *how you employ your perceptual capacity* to gain information via perceptual mechanisms. Yet, **1a** does not satisfy Siegel's criterion, since your perception differs due to a change of distal stimuli. Strictly speaking, your perceptual capacity and how it generates perceptual content are not influenced.

> **1b. Pine tree.** "Before and after a subject learns what pine trees look like, pine trees look different to her, and the visual experiences she has under the same external conditions differ in their content. But this is because gaining pine-tree-expertise makes her fixate on the shapes of the leaves on the trees. If a novice fixated the way the expert did, then her experience would have the same contents. The expertise influences experience content, by influencing fixation points." (*Ibid.*: 204–5)

1b is a case of perceptual learning, a long-term change in perceptual content that results from previous experiences (Connolly, 2017, 2019). Perceptual learning is sometimes called *diachronic cognitive penetration* (DCP; see also Burnston, 2021; Jenkin, 2023), as it involves the cognitive influence on perception. But **1b** also falls short of Siegel's criterion because experts pay more attention to the distinctive features that help identify pine trees than novices do while seeing the same distal stimuli. The background knowledge about pine trees makes an expert perceive more details even if the scene presented to her is identical to what is presented to a novice. In this case,

background knowledge does not influence the content formation process. What is influenced is the ability to process more specific stimuli. In other words, what is influenced in question is the "*conditions that determine what will get perceptually processed*" (Macpherson, 2012: 44), instead of the perceptual process *per se*. So, DCP does not satisfy Siegel's criterion because experts and novices do not attend to the same distal stimuli.

Therefore, only *synchronic cognitive penetration* (SCP) satisfies Siegel's criterion. Consider this case:

> **1c. Angry face.** "Jill believes, without justification, that Jack is angry at her. The epistemically appropriate attitude for Jill to take toward the proposition that Jack is angry at her is suspension of belief. But her attitude is epistemically inappropriate. When she sees Jack, her belief makes him look angry to her. If she didn't believe this, her experience wouldn't represent him as angry." (Siegel, 2012: 209)

1c is the case where occurrent cognitive states influence the perceptual process. Specifically, provided that the conditions that determine what will get perceptually processed are fixed (in this case, Jack's facial expression), the content of the penetrated experience would not be changed if the penetrating occurrent anger belief were not present.

There are reasons for the idea that SCP best explains **1c**. First, **1c** is not a case where cognition influences the way Jill employs perceptual capacities because the visual scene that appears to Jill does not change. Second, **1c** is not a case where the anger belief changes the conditions that determine what will get perceptually processed. Jill may experience Jack as angry or not angry by seeing the same facial expression. Attending to any feature on Jack's face does not suffice for the experience to represent anger. So, the formation of the anger experience cannot be explained by the change in attention or other perceptual factors. Following the two reasons above, the best explanation of **1c** is that the anger belief penetrates the perceptual process by altering the experience to represent anger.

We have seen the difference between **1a** (not a case of cognitive penetration), **1b** (DCP), and **1c** (SCP).[5] I do not take **1a** as a possible interpretation because no one would consider cases

[5] Different philosophers define cognitive penetration in different ways. Although Siegel precludes DCP from a case of cognitive penetration, Lyons's (2011) definition includes both DCP and SCP. But this is only a matter of labeling,

like **1a** as cognitive orientation. If what Chalmers means by cognitive orientation refers to cases like **1c**, then cognitive orientation is SCP. If he literally means that cognition influences perception, then cognitive orientation includes cases like **1b**, i.e., DCP. Section 6.2.2 develops the interpretation of SCP, and Section 6.2.3 addresses that of DCP. We will see that cognitive orientation can only be interpreted as DCP.

6.2.2. Cognitive Orientation as Synchronic Cognitive Penetration

Suppose what Chalmers means by cognitive orientation is SCP. The driver case is a case where sophisticated drivers perceive things in the rear mirror as located behind just because their background knowledge *synchronically* penetrates their perceptual process. Were the knowledge not present, the drivers, employing the same perceptual process, would perceive the things as located in the front. Likewise, in the VR case sophisticated users perceive virtual things as virtual because the background knowledge *synchronically* penetrates their perceptual process. Were the knowledge not present, the users, employing the same perceptual process, would perceive virtual things as physical. Hence, P2.2 is understood as an inference of the occurrence of SCP in the VR case from its occurrence in the driver case.

So, P2.2 is true because (1) the driver case is a case of cognitive orientation (SCP) because it satisfies the four features, and (2) given that the VR user case also satisfies the four features, it is another case of cognitive orientation (SCP). I assume (1) is true. However, as I will argue, satisfying the four features in the VR user case does not suffice for the occurrence of SCP in this case. P2.2 does not hold.

Since P2.2 is based on the thesis that satisfaction of the four features — *Knowledge*, *Familiarity*, *Action-dependence*, and *Naturalness* — indicates an occurrence of SCP, it will be largely weakened if a case satisfies the four features but does not involve SCP. This sub-section presents such a case, which I call the mirror-covering case, as it involves spatial inversion of individuals' visual experience by a set of mirrors covering their visual field. A paradigmatic mirror-covering case was shown in the Innsbruck Goggle Experiment on perceptual adaptation (Erismann, 1948; Kohler, 1951). In the experiment, the subjects were required to wear a set of mirror goggles through which they saw the world upside down for a few days. The goggles fully

as they both agree with the distinction of DCP and SCP. In this chapter I do not intend to adjudicate their disagreement. What I intend is to show the difference between DCP and SCP.

occupied the subjects' entire visual field. The participants were also required to do daily activities. During the first three days, the subjects could barely accomplish daily movements, such as grabbing or moving, because they were not used to the reversed visual scene. By the fifth day, the subjects' movement improved: "Things that had been seen upside down suddenly were upright once the participant brought their own hands in and traced the shapes they saw with their hands" (Kohler, 1951, as cited in Sachse et al., 2017). From the sixth day, the subjects started to act as if they had not been wearing the goggles. One subject could even draw a picture "in a quality as if drawn without wearing [the goggles]" (Sachse et al., 2017: 227).

The mirror-covering case satisfies the four features that indicate SCP. First, *Knowledge* certainly played a role in the subjects' experience in the mirror-covering case. The experimenters did not conceal the fact about the mirror goggles, and the subjects could recognize the presence of a reverse mirror while wearing the goggles. Also, the subjects became familiar with how the goggles functioned to their perception. For example, they were familiar with the fact that what appeared to them as located at the top was actually located at the bottom. To cope with the altered perception, the subjects eventually developed a new pattern of acting. In this way, *Familiarity* and *Action-dependence* are both satisfied. Lastly, *Naturalness* is also satisfied: The interpretation on which the perceived things are upside-down is more natural than the one on which they are as what they appear. That is, it was more natural for the subjects to interpret their visual experiences to be systematically altered.

However, SCP hardly occurs in the mirror-covering case. As reported, the subjects underwent a period of perceptual adaptation. For these subjects to cope with the environment, it could not be a sudden shift in what they perceived. The duration of the gradual adaptation period largely reduces the possibility of SCP occurrence. Since SCP requires the synchronic cognitive influence on perception, the subjects would not have spent a few days adapting to the altered perception if SCP had occurred in the mirror-covering case. In other words, if SCP were the case, then the subjects' knowledge that they were seeing the world upside-down would have penetrated their perception synchronically when they saw the world upside-down. So, they would have perceived the world similarly to those who did not wear the goggles. However, the subjects' inability to accomplish daily activities in the first few days shows that they did not perceive the upside-down world as we normally perceive it. This suggests that SCP hardly occurred in the

mirror-covering case. Although this case satisfies the four features of cognitive orientation, it is not a case of SCP.

Having the mirror-covering case on the table, it becomes unreasonable to infer the SCP occurrence from the driver case to the VR user case. Unlike the driver case, the VR user case resembles the mirror-covering case. The latter two cases do not involve discontinuity within the subjects' visual field. In the driver case, the rear mirror occupies a small proportion of the driver's visual field, so it is easy for them to visually compare the discontinuity of the mirror image to other visual regions. In this way, the perceptual content with respect to the mirror image is malleable to the content regarding the rest of the visual field. This influence could be SCP, as the content regarding other visual regions might be cognitive content that alters what the drivers perceive through the rear mirror. Yet, even if SCP occurs in the drive case, it is implausible to occur in the mirror-covering case or the VR user case. Since the mirror goggle or VR equipment occupies the entire visual field, no visual discontinuity is present in these cases. So, the only possible way for SCP to occur in these cases is the influence of the subject's background knowledge. However, as we have seen, the mirror-covering case does not involve the cognitive influence from background knowledge, suggesting that the subjects perceive what appears to them *as such*. Likewise, it is also untenable to suppose that SCP occurs in the VR user case because of the lack of visual discontinuity. Therefore, even if the four features of SCP are satisfied in the mirror-covering case and the VR user case, it is still implausible that SCP occurs in these cases.

Therefore, the VR user case is hardly a case of SCP. Although all the features that indicate SCP are satisfied by the VR user case and the mirror-covering case, it is still untenable to conclude that SCP occurs in these cases. Since the Innsbruck Goggle Experiment shows that SCP does not occur in the mirror-covering case, it is implausible that SCP would occur in the VR user case. Thus, it is not the case that sophisticated users perceive virtual things as virtual just because of SCP. If what Chalmers means by cognitive orientation is SCP, the analogical inference in P2.2 does not hold, rendering his argument unsound.

6.2.3. Cognitive Orientation as Diachronic Cognitive Penetration

Having rejected the interpretation of SCP, DCP is the remaining interpretation to make Chalmers's argument sound. This sub-section argues that DCP fits what Chalmers means by cognitive orientation, so it can support P2.2.

DCP is a long-term change in perceptual content that results from previous experiences/cognitive states (Connolly, 2017). It is typically considered as a sort of perceptual learning (Burnston, 2021; Jenkins, 2023). The learning process plays a significant role in DCP. Recall **1b**, where the perceivers' background knowledge enhances the sensitivity of their perceptual capacity. Perceivers knowing nothing about the tree category may perceive a pine tree merely *as a tree*. When they learn that the characteristics of pine trees can be found, say, through the needle-like shape of leaves, they start to fixate on such characteristics. Once they become adept at differentiating pine trees from others, they perceive a pine tree *as a pine tree*. The learning process relies on the knowledge of pine trees that changes how the perceivers perceive the tree. Echoing Macpherson's (2012) explanation, **1b** is the case where cognition influences the change from the condition that determines the content <tree> to the condition that determines the content <pine tree>.

The driver case can be interpreted in terms of DCP. Having learned that the rear mirror is present, naïve drivers become sophisticated and start to interpret the things in the mirror to be located behind. *Familiarity*, *Action-dependence*, and *Naturalness* may also feature in the drivers' learning process, such that the drivers become more familiar with how the rear mirror functions in their perception, how their actions depend on the mirror perception, and which interpretation (i.e., things located in the front vs. behind) is more natural to have. Through a learning process, the drivers no longer perceive the things *as located in the front*. Rather, they perceive the things in the mirror *as located behind*. In brief, their knowledge influences the change from the condition that determines the content <X being located in the front> to the condition that determines the content <X being located behind>.

P2.2 is plausible under the interpretation of DCP. The VR user case resembles the driver case in that both cases involve a learning process toward what is perceived. Naïve users become sophisticated by learning that they are interacting with virtual objects. As Chalmers (2017: 331) suggests, users' content does not change immediately when they learn the fact. Rather, they gradually become familiar with VR scenarios. Their action pattern also gradually turns out to depend on their knowledge of VR. Eventually, users' learning process changes their perception, altering the content associated with the reddish character from <X being non-virtually red> to <X being virtually red>. In this way, the change of content is plausibly made through VR users' learning process of what they perceive. Thus, the VR user case resembles the driver case if

cognitive orientation is understood with respect to DCP. The analogical inference of P2.2 is justified in this manner.

In sum, I have analyzed different interpretations of cognitive orientation and argued that only the interpretation of DCP helps justify P2.2. So, it is safely assumed that what Chalmers means by cognitive orientation is DCP. Once the interpretation is adopted, Chalmers succeeds in defending virtual veridicalism. The next section points out a further problem with Chalmers's defense.

6.3. Chalmers's Defense Revisited

Having the best version of Chalmers's argument for virtual veridicalism on the table, we are in a better position to assess its consistency with Chalmers's commitment to the five constraints of virtual veridicalism and Fregean representationalism (the Chalmers-Thompson account). This section points out two problems that hinder their consistency. The first problem regards the consistency between Fregean representationalism and the five constraints, especially the fifth one (*viz. Scope*). The second problem concerns the lack of resources in Fregean representationalism to explain the phenomenon of cognitive orientation (DCP).

At first glance, Chalmers's argument is sound when cognitive orientation is interpreted as DCP. As P2.2 is justified, C1 (i.e., sophisticated VR users perceive virtual objects as having virtual CLSs) follows. By C1 and P3 (i.e., virtual objects do have virtual CLSs), we have that sophisticated users perceive virtual objects as having the virtual properties they do instantiate. This statement implies C2: Many VR experiences are veridical.

However, an internal conflict arises if the argument is situated in the detailed picture of virtual veridicalism that Chalmers aims to defend. Specifically, Chalmers's argument is inconsistent with his commitment to the constraints of virtual veridicalism. C2 just is the claim of *Veridical VR Experience*. *Veridical Non-VR Experience* is satisfied because sophisticated VR users still veridically perceive physical objects as having physical properties. *Different Reliable Cause* is satisfied in that the virtual properties perceived by sophisticated users are categorically different from the physical properties represented by their ordinary experiences. Due to DCP, sophisticated users' reddish experiences differ in content within and outside VR. They perceive red objects as having the property of *virtual redness* in VR, but they perceive red objects as having the property of *physical redness* in ordinary environments. Assuming the phenomenal content view,

these experiences cannot be phenomenally identical because a difference in content entails a difference in phenomenal character: The reddish phenomenal character cannot determine both the content <physical redness> and the content <virtual redness>. Therefore, *Same Phenomenology* has to be rejected. If such an inconsistency is not resolved, Chalmers's defense of virtual veridicalism is significantly undermined.

There are two options to resolve the inconsistency. The first is to weaken Chalmers's commitment to virtual veridicalism by giving up some constraints. One might suggest that "reddish experiences" could be taken as a cluster of various phenomenal types, such that many phenomenally unidentical experiences may count as reddish experiences. Although both experiences of physical redness and those of virtual redness belong to reddish experiences, they need not be phenomenally identical. So, *Same Phenomenology* can be excluded from the constraints of virtual veridicalism.

I think this option, though possible, sacrifices too much. A motivation to assume the phenomenal identity between experiences of physical redness and those of virtual redness is that a subject cannot distinguish between them simply through what they appear to us. As I believe representationalists would accept this assumption, this solution would be the last option for them to resolve the inconsistency. Besides, if one insists that the subject must be able to distinguish between experiences of physical objects and those of virtual objects, one should carry the burden of proof. As I have argued in Chapter 3 that these experiences can be phenomenally identical, the objector should provide sufficient reasons to reject this point.

The second option, which I think is more reasonable, is to look for theoretical resources to resolve the inconsistency without giving up any virtual veridicalist constraint. In a nutshell, the inconsistency can be resolved if there is an account of content that not only grounds Chalmers's defense but also satisfies the relevant constraints, i.e., *Veridical VR Experience*, *Veridical Non-VR Experience*, *Same Phenomenology*, *Different Reliable Cause*, and *Scope*. Since Chalmers endorses Fregean representationalism, it is natural for him to apply this account to solve the inconsistency.

However, as I argued in Section 4.4.2, Fregean representationalism encounters a serious problem harmonizing *Scope* with other constraints. To avoid repetition, I briefly review this problem. This account claims that the phenomenal content is a Fregean proposition composed of modes of presentation (MoPs; Chalmers, 2004, 2006; Thompson, 2009). So, the phenomenal content of experience of a certain sort has a phenomenal MoP that picks out as extension *the*

property that normally causes experiences of that sort in me. This allows a VR experience and its corresponding non-VR experience[6] to pick out different properties as extension. Since reddish experience is normally caused by physical redness in ordinary environments but by virtual redness in virtual worlds, Fregean representationalism predicts that E_P represents physical redness and E_V represents virtual redness. In this way, *Same Phenomenology* and *Different Reliable Cause* are satisfied simultaneously. Moreover, both reddish VR experience and reddish non-VR experience are veridical because their objects do instantiate physical redness and virtual redness, respectively. *Veridical VR Experience* and *Veridical non-VR Experience* are satisfied.

However, Fregean representationalism does not satisfy *Scope*. Their consistency holds only in perfect virtual worlds (i.e., virtual worlds in which all entities that appear to the subjects are virtual entities) but not in mixed worlds (i.e., virtual worlds in which some entities that appear to the subjects are virtual entities and the others are physical entities). This account entails that phenomenal content picks out only one extension under a single circumstance. So, in a given circumstance, a reddish experience can represent either physical redness or virtual redness, but not both. In a mixed world, however, most reddish experiences (if not all) are jointly produced by physical redness and virtual redness. This suggests that most reddish experiences are falsidical in mixed worlds because some of their phenomenal content attributes wrong properties to the objects. As *Veridical VR Experience* is rejected in the mixed worlds, *Scope* is not satisfied by this account. Fregean representationalism is not an appropriate account of content in support of virtual veridicalism.

The second problem is that Chalmers's view on cognitive orientation (DCP) commits him to the contents of experiences that *change* over time, but his Fregean account lacks sufficient resources to explain such a change. Cognitive orientation (DCP) is a phenomenon in which cognition influences a change in perceptual content. During the process of cognitive orientation in a VR user, her reddish experience ceases to represent <physical redness> and then starts to represent <virtual redness>. Although Fregean representationalism can accommodate the veridicality of reddish experience before and after the content change, this account is unable to adequately explain how the change is associated with cognitive orientation (DCP).

[6] By definition (see footnote 1), a VR experience and its corresponding non-VR experience share the same phenomenal character.

According to Fregean representationalism, a normal cause is a statistically typical cause (see Section 4.4.1). What counts as the normal cause of a certain experience in a subject in a certain circumstance is what typically causes that experience in her throughout her perceptual history in that circumstance. So, her content change can be done by her mere exposure to the new circumstance for a certain period. In other words, the subject need not know anything about this new circumstance. She does not even need to know (or, in any sense, understand or grasp) that the CLSs she sees in this new circumstance are different from those that she saw in the previous circumstance. What is required for the content change is that she *in fact* be in the new circumstance for her experiences of CLSs to hook on the new typical cause.

However, allowing the subject not to know the change of content results in an epistemic gap for the subject to gain appropriate knowledge of the new circumstance. As the contents of CLSs before and after the change are supposedly phenomenal content, they are permanently associated with the relevant experiences. Yet, the change in this sort of content is not reflected in the relevant phenomenology. In particular, since experiences of *physical redness* and of *virtual redness* are both reddish experiences, there is no change in the reddish phenomenology during the change from the content <physical redness> to <virtual redness>. Through introspection, the subject is by no means aware of any difference in the experience. So, we may ask, by virtue of what can the subject get epistemic access to the content change? Presumably, a difference in content is supposed to be an epistemic difference. If the subject knows nothing different through the content change, it would be awkward to ascribe new content to her because she is unable to exploit the new content in further mental activities, e.g., forming a perceptual belief. Both *physical redness* and *virtual redness* appear red to the subject, so she cannot know which property is currently picked out. Thus, she may form a perceptual belief that something is red, but she cannot form a belief that something is physically red or virtually red. She may also be able to believe that the objects she saw before and after change are all red, but she cannot form a belief that these objects have the same or different properties. Consequently, the subject has no epistemic access to the property that is predicted to be represented in her experience. Due to this epistemic gap, it would be awkward to ascribe contents like <physical redness> or <virtual redness> to the subject's reddish experience even if she has been exposed to the appropriate circumstance for a period of time.

This epistemic gap substantiates the importance of the subject's background knowledge of the circumstance she is in. As Chalmers emphasizes the process of cognitive orientation (DCP) during the content change:

> "A naïve user who does not *know* they are using virtual reality will undergo the illusion that certain objects are present in physical space in from of them. After they *learn* they are using virtual reality, the perceptual illusion may persist for a period, but they will not be fooled into believing that the objects are present. After some time, a sophisticated user will become *familiar* with VR, and they will act in ways that turn on *interpreting* themselves to be in VR. […] When the sophisticated user of VR *knows* they are looking at virtual objects, they have a distinctive *phenomenology of virtuality*. […] When a sophisticated user has the phenomenology of virtuality, it is plausible that they *perceive* the objects they are interacting with as being virtual objects in virtual space." (Chalmers, 2017: 331–2; italics mine)

This passage points out the importance of the subject's background knowledge to the content change. The gist is that the change of content occurs only when the subject *knows* that she is in a new circumstance whose environmental factors are significantly different from those in the previous circumstance. Because of her background knowledge, the subject has a distinctive phenomenology of virtuality that makes her perceptual contents in a virtual world different from those in an ordinary environment. Therefore, cognitive orientation (DCP) requires background knowledge, but Fregean representationalism lacks sufficient resources to explain the crucial role that background knowledge plays in content change. In order for Chalmers's defense of virtual veridicalism to be complete, he would need an extra explanation other than Fregean representationalism to fill this gap.

To sum up, apart from the inconsistency between Fregean representationalism and the five virtual veridicalist constraints, this account does not have sufficient resources to explain cognitive orientation (DCP), which is central to Chalmers's argument for virtual veridicalism. If there is another account of content that not only coheres with the five constraints but also sufficiently explains cognitive orientation (DCP), that account is a better ground for Chalmers's defense of

virtual veridicalism than Fregean representationalism does. The next section turns to role representationalism and assesses whether it is a better account.

6.4. Role Representationalism to the Rescue

After pointing out the problems for Fregean representationalism to cohere with Chalmers's defense of virtual veridicalism, this section turns to role representationalism and examines if it can avoid these problems and appropriately ground Chalmers's defense. I argue that role representationalism works consistently with all the virtual veridicalist constraints and accommodates the phenomenon of cognitive orientation (DCP). Therefore, role representationalism is the best account for Chalmers to presuppose in order to defend virtual veridicalism. Again, since role representationalism is extensively defended in Chapter 5, this section briefly summarizes the account and focuses on how this view serves as an appropriate ground for Chalmers's defense.

Role representationalism combines representationalism about perceptual experience and role functionalism of lenient content (Section 5.2):

> *Role Functionalism about Lenient Content*: The lenient content of a perceptual experience with phenomenal character *C* is a proposition of the form <something is *F*>, where *F* is the role property of *having a realizer property that normally causes experiences with C in the subject under the current perceptual circumstance*.

As argued in Section 5.3, role representationalism is coherent with all the virtual veridicalist constraints that Chalmers is committed to. First, it accommodates *Same Phenomenology* and *Different Reliable Cause* at once. The lenient contents of E_P and E_V both attribute the role property of *having a realizer property that normally causes reddish experiences in me under the current perceptual circumstance*. The role property happens to be realized by distinct realizers in different circumstances. In an ordinary environment, physical redness normally causes reddish experiences, so it realizes the role property. In a virtual world, in contrast, virtual redness realizes the same role property also because it normally causes reddish experiences. Thus, although E_P and E_V are reliably caused by distinct properties, they share the same phenomenology. Second, given that the property represented by reddish experiences is instantiated in both virtual and non-virtual circumstances, it

is plausible for both E_P and E_V to be veridical. *Veridical VR Experience* and *Veridical Non-VR Experience* are both satisfied.

Lastly, this account satisfies *Scope*. Suppose that E_V occurs in a mixed world instead of a perfect virtual world. The change of E_V's circumstance does not affect the satisfaction of *Same Phenomenology, Different Reliable Cause*, and *Veridical Non-VR Experience*. E_V still has the same phenomenology as E_P does. E_V is still caused by VC_{600}, which is categorically different from what causes E_P, *viz*. C_{600}. The change in E_V's circumstance does not affect the veridicality of E_P. In addition, *Veridical VR Experience* is satisfied. Role representationalism predicts that E_V's lenient content attributes the role property of *having a realizer property that normally causes reddish experience in me under the current perceptual circumstance*. As VC_{600} is a typical cause of reddish experience in the mixed world, it is natural for the subject, with sufficient knowledge, to expect VC_{600} to cause her reddish experience in the mixed world; such an expectation qualifies VC_{600} as the normal cause under this circumstance. Since the instantiation of VC_{600} necessitates the instantiation of the role property attributed by E_V's lenient content, E_V is leniently veridical in the mixed world in that its object instantiates the property that E_V's lenient content attributes. Thus, *Scope* is satisfied. Role representationalism contains sufficient resources to accommodate Chalmers's commitment to virtual veridicalism.

Not only does role representationalism cohere with Chalmers's commitment to virtual veridicalism, but it is also consistent with Chalmers's argument from cognitive orientation (DCP). Recall that cognitive orientation is perceptual learning. VR users who initially perceived virtual objects as having physical properties learn to perceive them as having virtual properties. Supposedly, both the pre-learning experience and the post-learning experience are reddish experiences, so they share the same phenomenology. However, as I have pointed out in the previous section, the content change during the process is the core that makes Chalmers's argument conflict with his commitment to *Same Phenomenology*. If the contents of the pre-learning experience and the post-learning experience are different, and if the contents are phenomenal contents, their phenomenology must differ. So, Chalmers's argument seems inconsistent with his commitment to virtual veridicalism. This is why we seek an appropriate account of phenomenal content to resolve the inconsistency.

Role representationalism can help resolve the inconsistency. In a nutshell, the content that changes through perceptual learning is not lenient content. That is, <physical redness> and <virtual

redness> are contents that are hidden from phenomenology. In this way, role representationalism does not encounter the epistemic gap faced by Fregean representationalism. *Physical redness* and *virtual redness* are the normal causes of reddish experiences in their respective circumstances. But they are not the property attributed by the lenient content of reddish experiences. Rather, what is attributed is the property realized by physical redness and virtual redness in these circumstances. Put differently, cognitive orientation (DCP) does not affect the lenient content. The learning process merely affects what particular property realizes the attributed property. Initially, the lenient content of reddish experience attributes the role property that is realized by physical redness in ordinary environments. When the subject perceives something red in VR after the learning process, the lenient content of her reddish experience still attributes the same property, but the property is realized by a different property, viz. virtual redness, because she takes virtual redness as a normal cause in the new circumstance by learning what would happen in VR. Therefore, while being reliably caused by distinct properties, the pre-learning experience and the post-learning experience still share the same phenomenology. This not only preserves our intuition that the pre-learning and post-learning experiences do not yield an epistemic difference but also explains the difference of (phenomenally inert) contents involved in these experiences. Thus, *Same Phenomenology* no longer conflicts with the process of cognitive orientation (DCP). Role representationalism is resourceful in removing the inconsistency between Chalmers's defense and his commitment to virtual veridicalism.

In addition, role representationalism plausibly accommodates Chalmers's argument from cognitive orientation (DCP). As discussed in Section 5.2.4, role representationalism takes what counts as "normal" for a subject in a certain circumstance as equal to what the subject expects to be in that circumstance, and the subject's expectation depends on her background knowledge about the circumstance. So, it is easy to see that role representationalism allows for content change via learning because the content is derived from what we expect to attribute to our experiences. Specifically, learning from our background knowledge, we expect our experience in the circumstance to be caused by its typical cause. As a result, an erudite subject will no longer expect physical redness to be what will cause her reddish experience in a virtual world. Instead, she starts to expect virtual redness to be the cause. Although such a change does not affect the lenient content, it affects what in fact plays the role of normally causing the experience. Hence, the process of

cognitive orientation (DCP) is a process whereby the subject expects a new property to cause her experience in a new circumstance.

In sum, role representationalism is the most appropriate account of content that assuages the conflict between Chalmers's argument from cognitive orientation (DCP) and his commitment to virtual veridicalism. Chalmers's defense of virtual veridicalism cannot succeed unless he adopts an account of content to resolve the conflict. Instead of Fregean representation, the account endorsed by Chalmers himself, role representationalism contains sufficient resources to harmonize the conflict. Therefore, for Chalmers's argument to succeed in defending virtual veridicalism, role representationalism has to be presupposed as the underlying theory of content.

6.5. Conclusion

This chapter assessed Chalmers's defense of virtual veridicalism. I first presented and disambiguated Chalmers's argument, arguing that its central notion, "cognitive orientation", should be interpreted as a sort of diachronic cognitive penetration. Interpreting this way, I argued that an internal inconsistency takes place when combining Chalmers's argument with the view he aims to defend. I also suggested that an appropriate account of content is required to underpin his argument so as to resolve the inconsistency. Furthermore, role representationalism is the best account that grounds Chalmers's defense. Therefore, I conclude that Chalmers needs to presuppose role representationalism for his argument to succeed in defending virtual veridicalism. As a result, Chalmers's defense of virtual veridicalism is incomplete without my own defense of the view.

Chapter 7: Conclusion

7.1. On What I Have Done Thus Far

In the preceding chapters, my primary concern has been to comprehensively defend virtual veridicalism, according to which many VR experiences are leniently veridical. I believe my defense positions me to answer the questions raised in Chapter 1:

1) *What is there in virtual reality*?
2) *How do we perceive virtual things*?

The first question concerns the metaphysics of VR. It was addressed in Chapter 2: There are virtual objects and virtual properties in VR. This chapter offered a detailed picture of the nature of these entities. I defended *weak virtual digitalism*, the view that virtual objects exist and are metaphysically dependent upon distinct, existing data structures. For the instantiation of virtual properties, I distinguished between two kinds of virtual properties. *Categorical properties* that are instantiated by virtual objects are not instantiated by non-virtual ordinary objects. For example, a virtual kitten is never a kitten, since *being a virtual kitten* is categorically different from *being a kitten*. In contrast, *role properties* that are instantiated by virtual objects may also be instantiated by non-virtual ordinary objects. For instance, a virtual library is a library because both virtual and non-virtual libraries instantiate the same property, namely, *being a library*. As virtual objects instantiate categorical and role properties, both virtual objects and virtual properties are real in VR. This realist view serves as a foundation for my arguments for virtual veridicalism. I believe it is also sufficient for answering the first question.

Chapters 3 – 5 addressed the second question. Defending virtual veridicalism, these chapters provided arguments from different perspectives — i.e., from the veridicality of experience and the nature of content. Chapter 3 argued for virtual veridicalism by re-examining our conception of perceptual veridicality. I distinguished between two criteria of veridicality. The *strict criterion* sets an overly high standard to count an experience as veridical. While the criterion is widely employed by philosophers of perception, it is not what common people mean by a perceptual experience being "veridical". In contrast, the *lenient criterion* fits common sense better, as we usually consider our ordinary experiences veridical on the grounds that they function well for our

practical purposes. Given the lenient criterion, we should also admit that many VR experiences are veridical. So, from the perspective of veridicality, virtual veridicalism is a plausible view.

Chapters 4 and 5 looked for a picture of the lenient content of perceptual experience that coheres with virtual veridicalism. Chapter 4 searched for the appropriate account within the representationalist camp. Evaluating three currently prominent accounts — physical Russellian representationalism, dispositionalist Russellian representationalism, and Fregean representationalism — I argued that none of them provide a coherent picture of perceptual experience with what virtual veridicalism requires. Instead, Chapter 5 proposed role representationalism, according to which an experience leniently represents the role property of *having a realizer property that normally causes the experience of that sort in the subject under the current perceptual circumstance*. I also argued that role representationalism provides sufficient support to virtual veridicalism with a coherent picture of the nature of perceptual content. Therefore, the second question can be clearly answered: We perceive virtual objects and properties leniently veridically.

Chapter 6 examined Chalmers's (2017) defense of virtual veridicalism. I argued that his defense is insufficient unless it accepts role representationalism as an underlying account of the nature of perceptual content. However, it is also possible to integrate my defense with his argument from cognitive orientation (interpreted as diachronic cognitive penetration). As virtual worlds are practically subordinate to the physical world, our perception is supposed to anchor ordinary objects and properties rather than virtual ones. Although virtual veridicalism is justified in Chapters 3 – 5, the view would become more comprehensive if it could explain *how VR experiences become veridical*. I believe Chalmers's argument sheds some light on the causal process in which a naïve VR user *learns* to perceive virtual objects and properties veridically. Incorporating cognitive orientation, virtual veridicalism is not only theoretically plausible but also practically probable.

7.2. On What We Can Do in the Future

The preceding chapters provided epistemological and metaphysical foundations for the study of VR, but many relevant issues remain untouched in this dissertation. Before ending this dissertation, I want to highlight a few issues worth further investigating in the future. Firstly, the meaning of life in VR is widely debated. Nozick (1974), in his famous thought experiment on the experience machine, once asked if we would choose to live in a well-programmed VR for the rest

of our life. Chalmers also addressed a similar question in his book Reality+ (2022), asking whether we can lead a good life in a virtual world. Although they have different attitudes toward the answers, these questions suggest the importance of probing the meaning of life in the new technology. Although this dissertation is silent on this issue, its conclusions may shed light on the relevant debates. For instance, given the reality of virtual entities and the veridicality of VR experiences, VR seems to be in a good position to embody values for our lives. If it is possible to have a virtual family that could assuage loneliness, ignite hopes, and share emotions, living a meaningful life in VR might not be an absurd choice.

Secondly, the positive conclusions about the epistemological and metaphysical foundations of VR seem to urge us to enact a series of rules regulating people's conduct in the virtual realm. As addressed in Chapter 1, Zuckerberg's vision of building a distance-free society in the Metaverse becomes plausible if these foundations are shared by its members. I believe this dissertation provided sufficient grounds for constructing societies in the virtual world. Specifically, we know that something in VR is real and can be veridically perceived. We can further form regulations on these real entities based on moral values or social consensus. For instance, there could be a regulation on virtual theft (Wildman & McDonnell, 2020), virtual murder (Luck, 2009), and equality in VR (Frank, 2017). Of course, there is a long way to make the virtual society cohesive. I believe my discussion of these foundational issues can contribute to our progress toward this ultimate goal.

Thirdly, VR not only creates a new area for philosophical investigation but also serves as a wonderful tool for us to probe our own nature. VR has been applied to research in areas such as education, psychiatry, and social psychology. Moreover, it is employed as an advanced tool to investigate the nature of self-consciousness. Since VR creates imaginative scenarios that can induce experiences that can hardly take place in ordinary environments, researchers can distinguish between the fundamental and contingent aspects of human self-consciousness more easily through experimental manipulations. For example, some work I have co-authored uses VR to demonstrate that some long-held assumptions might be wrong: Liou et al. (2024) argues that the first-person perspective is not as fundamental for embodied self-consciousness as we commonly assume; Liang et al. (2018) contends that an illusory sense of agency plus an appropriate manipulation of the first-person perspective may augment the flexibility of the sense of body ownership, such that one may feel as if one has four hands simultaneously; Huang et al. (2017) shows that the sense of self-

location is influenced by the integration of sensory information received from the first-person perspective and from the third-person perspective. Thanks to VR, our investigations can delve into realms we have never explored.

Finally, I envision a future where philosophy and technology are tightly integrated. Recall the method of *technophilosophy*. This dissertation asks philosophical questions about VR and uses VR to help answer traditional philosophical questions. This reciprocal relationship between philosophical progress and technological development can be applied to other technologies. For instance, in Section 4.4.3, we have seen the theoretical connection between VR and MR or AR through the problem of mixed reality. If this problem is pushed one step forward, the technology of holograms could also be taken into consideration, as it is said to create a virtual world without the aid of concrete VR devices. We can also ask philosophical questions about holographic technology and use holographic technology to answer philosophical questions with a new solution. To end this dissertation, I echo the statement that ended the introduction of Chapter 1: Technology finds its place in the modern world with the support of philosophical foundations while simultaneously providing new resources to reframe longstanding philosophical debates.

Bibliography

Akins, K. (1996). Of sensory systems and the "aboutness" of mental states. *Journal of Philosophy* ***93***(7): 337–372.

Audi, P. (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy* ***109***(12): 685 –711.

Baysan, U. (2015). Realization Relations in Metaphysics. *Minds & Machines* ***25***: 247–260.

Beisbart, C. (2019). Virtual Realism: Really Realism or only Virtually so? A Comment on D. J. Chalmers's Petrus Hispanus Lectures. *Disputatio* ***11***(55): 297–331.

Bennett, K. (2007). Mental Causation. Philosophy. *Philosophy Compass* ***2***(2): 316–337.

Berkeley, G. (1713/1979). *Three Dialogues between Hylas and Philonous.* Ed. Robert M. Adams. Indianapolis: Hackett.

Berkeley G. (1734/1982). *A Treatise Concerning the Principles of Human Knowledge*. Hackett, New York.

Bird, A. (2007). *Natures Metaphysics: Laws and properties*. Oxford: Oxford University Press.

Block, N. (1990). Inverted Earth. *Philosophical Perspectives* ***4***: 53–79.

Block, N. (1996). Mental Paint and Mental Latex. *Philosophical Issues* ***7***: 19–49.

Boghossian, P. A., Velleman, J. D. (1989). Color as a secondary quality. *Mind* ***98***: 81–103.

Bostrom, N. (2003). Are We Living in a Computer Simulation? *Philosophical Quarterly* ***53***(211): 243–255.

Bourget, D. (2017). Representationalism and Sensory Modalities: An Argument for Intermodal Representationalism. *American Philosophical Quarterly* ***54***(3): 251–268.

Bourget, D. & Mendelovici, A. (2014). Tracking Representationalism. In Andrew Bailey (ed.), *Philosophy of Mind: The Key Thinkers* (pp. 209-235). Continuum.

Burnston, D. C. (2021). Perceptual Learning, Categorical Perception, and Cognitive Permeation. *Dialectica* ***75***(1): 25–58.

Byrne, A. (2020). Inverted Qualia. *Stanford Encyclopedia of Philosophy.* In E. N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2020/entries/qualia-inverted/>

Byrne, A. & Hilbert, D. R. (1997). Colors and reflectances. In A. Byrne & D. R. Hilbert (eds.), *Readings on Color, Volume 1: The Philosophy of Color* (Ch. 14). MIT Press.

Byrne, A. & Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences* ***26***(1): 3–21.

Carnap, R. (1967). *The Logical Structure of the World: And, Pseudoproblems in Philosophy*, trans. Rolf A. George. Berkeley: University of California Press.

Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Chalmers, D. J. (2002). On sense and intension. *Philosophical Perspectives* ***16***: 135–82.

Chalmers, D. J. (2004). The representational character of experience. In Brian Leiter (ed.), *The Future for Philosophy* (pp. 153–181). Oxford University Press.

Chalmers, D. J. (2006). Perception and the fall from Eden. In T. S. Gendler & J. Hawthorne (eds.), *Perceptual Experience* (pp. 49-125). Oxford University Press.

Chalmers, D. J. (2012). *Constructing the World*. Oxford: Oxford University Press.

Chalmers, D. J. (2017). The Virtual and the Real. *Disputatio* **9**(46): 309–352.

Chalmers, D. J. (2018). Structuralism as a Response to Skepticism. *Journal of Philosophy* ***115***(12): 625–660.

Chalmers, D. J. (2019). The Virtual as the Digital. *Disputatio* ***11***(55): 453–486.

Chalmers, D. J. (2021). Inferentialism, Australian style. Proceedings and Addresses of the *American Philosophical Association* ***92***.

Chalmers, D. J. (2022). *Reality+: Virtual worlds and the problems of philosophy*. W. W. Norton & Company.

Chilovi, S. (2021). Grounding entails supervenience. *Synthese* ***198***(Suppl 6): 1317–1334.

Choi, S. & Fara, M. (2021). Dispositions. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/spr2021/entries/dispositions/>.

Cohen, J. (2005). Colors, functions, realizers, and roles. *Philosophical Topics* ***33***(1): 117–140.

Cohen, J. (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford, GB: Oxford University Press UK.

Connolly, K. (2017). Perceptual Learning. In Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/sum2017/entries/perceptual-learning/>.

Connolly, K. (2019). *Perceptual Learning: The Flexibility of the Senses*. Oxford: Oxford University Press.

Cutter, B. (2018). Paradise Regained: A Non-Reductive Realist Account of the Sensible Qualities. *Australasian Journal of Philosophy* ***96***(1): 38–52.

Cutter, B. (2021). Perceptual illusionism. *Analytic Philosophy* ***62***(4): 396–417.

Dasgupta, S. (2014). On the Plurality of Grounds. *Philosophers' Imprint* ***14***(20): 1–28.

Dicker, G. (2011). *Berkeley's Idealism: A Critical Examination*. USA: Oxford University Press.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.

Egan, A. (2006a). Appearance properties? *Noûs* ***40***(3): 495–521.

Egan, A. (2006b). Secondary Qualities and Self-Location. *Philosophy and Phenomenological Research* ***72***(1): 97–119.

Elder-Vass, D. (2010). *The Causal Power of Social Structures: Emergence, Structure and Agency.* Cambridge University Press.

Erismann, T. (1948). Das Werden der Wahrnehmung. Die Gesichtswahrnehmung in ihrer Abhangigkeit von anderen Sinnesgebieten und von der eigenen Vergangenheit. In *Kongressbericht; Berufsverband Deutscher Psychologen, Bonn 29. August bis 2. September 1947* (pp. 51-86). Hamburg: Nolke.

Fodor, J. (1974). Special sciences. *Synthese* ***28***: 97–115.

Foster, J. (1994). In Defence of Phenomenalistic Idealism. *Philosophy and Phenomenological Research* ***54***(3): 509–529.

Foster, J. (2008). *A world for us: the case for phenomenalistic idealism*. New York: Oxford University Press.

Frank, M. A. (2017). The Desert of the Unreal: Inequality in Virtual and Augmented Reality. *UC Davis Law Review 51, University of Miami Legal Studies Research Paper No. 17-24*.

Frege, G. (1892). Uber Sinn und Bedeudung. *Zeitschrift für Philosophie und philosophische Kritik* **100**, 25-50. Translated as "On Sense and Reference" in Geach and Black (eds. and trans.), 1980, pp.56–78.

Gillett, C. (2002). The dimensions of realization: A critique of the standard view. *Analysis* ***62***: 316–323.

Gillett, C. (2003). The metaphysics of realization, multiple realizability and the special sciences. *Journal of Philosophy* ***100***: 591–603.

Gillett, C. (2010). Moving beyond the subset model of realization: The problem of qualitative distinctness in the metaphysics of science. *Synthese* ***177***: 165–192.

Grabarczyk, P. (2024). What is Real in Virtual Reality? *Roczniki Filozoficzne* ***72***(1): 79–98.

Green, E. J. & Rabin, G. O. (2020). Use Your Illusion: Spatial Functionalism, Vision Science, and the Case Against Global Skepticism. *Analytic Philosophy* ***61*** (4): 345–378.

Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Hackett.

Heil, J. (2003). *From an Ontological Point of View*. Oxford: Oxford University Press.

Heil, J. (2005). Dispositions. *Synthese* ***144***: 343–356.

Helmholtz, H. L. F. von (1877/1954). *On the sensations of tone as the physiological basis for the theory of music* (2nd ed.; A. J. Ellis, Trans.). New York: Dover.

Huang, H.-C., Lee, Y.-T., Chen, W.-Y., & Liang, C. (2017). The Sense of 1PP-Location Contributes to Shaping the Perceived Self-location Together with the Sense of Body-Location. *Frontiers in Psychology* ***8***(370): 1–12.

Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield.

Hurley, S. L. (1998). *Consciousness in Action*. Harvard University Press.

Ishii, T. R., & Atkins, P. Essential vs. Accidental Properties. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/spr2023/entries/essential-accidental/>.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. New York: Oxford University Press.

Jenkin, Z. (2023). Perceptual learning and reasons-responsiveness. *Noûs* ***57***(2): 481–508.

Kenyon, K., Kinakh, V., & Harrison, J. (2023). Social virtual reality helps to reduce feelings of loneliness and social anxiety during the Covid-19 pandemic. *Sci Rep* ***13***: 19282.

Kim, J. (1993). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge University Press, New York.

Kim, J. (1998). *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: Bradford.

Kim, J. (2005). *Physicalism, or Something Near Enough*. Princeton, NJ: Princeton University Press.

Kohler, I. (1951). *Uber Aufbau und Wandlungen der Wahrnehmungswelt. Insbesondere uber 'bedingte Empfindungen'*. Wien: Rohrer: Osterreichische Akademie der Wissenschaften, Philosophisch-historische Klasse. Sitzungsberichte 227, Band 1.

Kriegel, U. (2023). The Structure of Phenomenal Justification. *Australasian Journal of Philosophy* ***101***(2): 282–297.

Laffitte, A., Neiers, F. & Briand, L. (2017). Characterization of taste compounds: chemical structures and sensory properties. In E. Guichard, C.Salles, M. Morzel, A.-M. Le Bon (eds.), *Flavour: From Food to Perception* (Ch. 7). Wiley & Sons, Ltd.

Langer, S. K. (1953). *Feeling and Form. A Theory of Art.* Charles Scribner's Sons: New York.

Lawson, T. (2012). Ontology and the study of social reality: emergence, organisation, community, power, social relations, corporations, artefacts and money. *Cambridge Journal of Economics* ***36***(2): 345–385.

Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* ***50***: 249–258.

Lewis, D. (1979). Attitudes de dicto and de se. *Philosophical Review* ***88*** (4): 513–543.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Lewis, D. (1994). Reduction of Mind. In S. Guttenplan, (ed.), *A Companion to Philosophy of Mind,* (pp. 412–31). Oxford: Blackwell,

Liang, C., Lee, Y.-T., Chen, W.-Y., & Huang, H.-C. (2018). Body-as-Subject in the Four-Hand Illusion. *Frontiers in Psychology* ***9***(1710): 1–9.

Liou, W.-K., Lin, W.-H., Lee, Y.-T., Chen, S., & Liang, C. (2024). The distinction between first-person perspective and third-person perspective in virtual bodily self-consciousness. *Virtual Reality* ***28***: 1.

List, C., & Menzies, P. (2009). Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy* **106**(9): 475–502.

Locke, J. (1689/1975). *An essay concerning human understanding.* J. W. Yolton (ed.) Oxford: Oxford University Press.

Luck, M. (2009). The gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology* ***11***(1): 31–36.

Ludlow, P. (2019). The Social Furniture of Virtual Worlds. *Disputatio* ***11***(55): 345–369.

Lycan, W. G. (1996). *Consciousness and Experience*. MIT Press.

Lyons, J. (2011). Circularity, reliability, and the cognitive penetrability of perception. *Philosophical Issues* ***21***(1): 289–311.

Macpherson, F. (2012). Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism. *Philosophy and Phenomenological Research* ***84***(1): 24–62.

Margot, C. (2009). A noseful of objects. *Nature neuroscience* ***12***(7): 813–814.

Maselli, A. & Slater, M. (2013). The building blocks of the full body ownership illusion. *Front. Hum. Neurosci.* ***7***: 83.

Maurin, A. (2019). Grounding and Metaphysical Explanation: It's Complicated. *Philosophical Studies* ***176***(6): 1573 –1594.

McDonnell, N. & Wildman, N. (2019). Virtual Reality: Digital or Fictional? *Disputatio* ***11***(55): 371–397.

McLaughlin, B. (2003). The place of color in nature. In R. Mausfeld and D. Heyer (eds.), *Colour Perception: Mind and the Physical World,* (pp. 475–502). Oxford University Press: New York.

Mendelovici, A. (2013). Reliable Misrepresentation and Tracking Theories of Mental Representation. *Philosophical Studies* ***165***(2): 421–443.

Mendelovici, A. (2018). *The Phenomenal Basis of Intentionality*. New York, USA: Oxford University Press.

Mendelovici, A. (2019). Immediate and Reflective Senses. In S. Gouveia, M. Curado & D. Shottenkirk (eds.), *Perception, Cognition and Aesthetics* (pp. 187–209). New York: Routledge Studies in Contemporary Philosophy.

Menzies, P., & List, C. (2010). The Causal Autonomy of the Special Sciences. In Cynthia McDonald & Graham McDonald (eds.), *Emergence in Mind*. Oxford, U.K.: Oxford University Press. pp. 108–129.

Metzinger, T. (2018). Why Is Virtual Reality Interesting for Philosophers? *Front. Robot. AI* ***5***: 101.

Milgram, P., & Kishino, F. (1994). A taxonomy of mixed reality visual displays. *IEICE Trans. Inform. Syst*. **77**: 1321–1329.

Millar, B. (2013). Colour constancy and Fregean representationalism. *Philosophical Studies* ***164***(1): 219–231.

Millikan, R. G. (1984*). Language, Thought, and Other Biological Categories: New Foundations for Realism*. MIT Press.

Mumford, S. (1998). *Dispositions*. Oxford: Oxford University Press.

Mumford, S. (2004). *Laws in Nature*. London: Routledge.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review* ***83***: 435–50.

Neander, K. (2017). *A Mark of the Mental: A Defence of Informational Teleosemantics*. Cambridge, MA: The MIT Press.

Nida-Rümelin, M. (1996). Pseudonormal Vision: An Actual Case of Qualia Inversion? *Philosophical Studies* ***82***(2):145–157.

Nozick, R. (1974). *Anarchy, State, and Utopia*. New York: Basic Books.

Orilia, F. & Paoletti, M. P. (2022). Properties. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/spr2022/entries/properties/>.

Pautz, A. (2006). Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Noûs* ***40***(2): 205–240.

Pautz, A. (2014). The Real Trouble with Armchair Arguments against Phenomenal Externalism. In M. Sprevak, J. Kallestrup (eds.), *New Waves in Philosophy of Mind.* Palgrave Macmillan, London.

Pautz, A. (2021). Consciousness meets Lewisian interpretation theory: A multistage account of intentionality. In Uriah Kriegel (ed.), *Oxford Studies in Philosophy of Mind*, Vol. 1.

Peacocke, C. (1983). *Sense and Content: Experience, Thought and Their Relations*. Oxford University Press.

Peacocke, C. (1992). *A Study of Concepts*. MIT Press.

Pereboom, D. (2013). Russellian Monism and Absolutely Intrinsic Properties. In Uriah Kriegel (ed.), *Current Controversies in Philosophy of Mind* (pp. 40–69). Routledge.

Perry, J. (1979). The problem of the essential indexical. *Noûs* ***13***(1): 3–21.

Polger, T. (2004). *Natural minds*. Cambridge, MA: MIT Press.

Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs* ***34***(4): 517–49.

Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (eds.), *Art, mind and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.

Ramsey, F. P. (1954). Theories. In R. B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*. London: Routledge and Kegan Paul.

Robb, D., Heil, J., & Gibb, S. (2023). Mental Causation. In Edward N. Zalta & Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. URL = <https://plato.stanford.edu/archives/spr2023/entries/mental-causation/>.

Rupert, R. D. (2006). Functionalism, mental causation, and the problem of metaphysically necessary effects. *Noûs* ***40***(2): 256–83.

Sachse, P., Beermann, U., Martini, M., Maran, T., Domeier, M., & Furtner, M. R. (2017). "The world is upside down" – The Innsbruck Goggle Experiments of Theodor Erismann (1883–1961) and Ivo Kohler (1915–1985). *Cortex* **92**: 222–232.

Schaffer, J. (2005). Quiddistic Knowledge. *Philosophical Studies* ***123***: 1–32.

Schouten, J. F. (1940). The Residue, a New Concept in Subjective Sound Analysis. *Proceedings of the Koninklijke Nederlandse Akadademie* ***43***: 356–65.

Segal, G. M. A. (2009). The Causal Inefficacy of Content. *Mind and Language* ***24***(1): 80–102.

Shockley, K., Carello, C. & Turvey, M. T. (2004). Metamers in the haptic perception of heaviness and moveableness. *Perception & Psychophysics* ***66***: 731–742.

Shoemaker, S. (1982). The inverted spectrum. *Journal of Philosophy* ***79***: 357–381.

Shoemaker, S. (1994). Self-knowledge and "Inner-sense". *Philosophy and Phenomenological Research* ***54***: 249–314.

Shoemaker, S. (1996). *The First Person Perspective and Other Essays*. Cambridge University Press.

Shoemaker, S. (2000). Phenomenal Character Revisited. *Philosophy and Phenomenological Research* ***60***: 465–467.

Shoemaker, S. (2001a). Realization and mental causation. In C. Gillett & B. Loewer (eds.), *Physicalism and its discontents* (pp. 74–98). Cambridge: Cambridge University Press.

Shoemaker, S. (2001b). Introspection and Phenomenal Character. *Philosophical Topics* ***28***(2): 247–273.

Shoemaker, S. (2003). Content, Character and Color. *Philosophical Issues* ***13***: 253–278.

Shoemaker, S. (2006). On the Way Things Appear. In T. S. Gendler & J. Hawthorne (eds.), *Perceptual Experience* (pp. 461–480). Oxford University Press.

Shoemaker, S. (2007). *Physical realization*. Oxford: Oxford University Press.

Shoemaker, S. (2013). Physical realization without preemption. In S. C. Gibb, E. J. Lowe, & R. D. Ingthorsson (eds.), *Mental causation and ontology* (pp. 35–57). Oxford: Oxford University Press.

Siegel, S. (2012). Cognitive Penetrability and Perceptual Justification. *Noûs* **46**(2).

Skarbez, R., Smith, M., & Whitton, M. C. (2021). Revisiting Milgram and Kishino's Reality-Virtuality Continuum. *Front. Virtual Real.* **2**:647997.

Slater, M. (2009). Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* ***364***(1535): 3549–3557.

Slater, M. & Sanchez-Vives, M. V. (2016). Enhancing Our Lives with Immersive Virtual Reality. *Front. Robot. AI* ***3***: 74.

Slater, M. & Sanchez-Vives, M. V. (2022). Is Consciousness First in Virtual Reality? *Front. Psychol.* ***13***: 787523.

Thompson, B. (2007). Shoemaker on phenomenal content. *Philosophical Studies* ***135***(3): 307–334.

Thompson, B. (2009). Senses for senses. *Australasian Journal of Philosophy* ***87***(1): 99–117.

Thompson, B. (2010). The Spatial Content of Experience. *Philosophy and Phenomenological Research* ***81***(1): 146–184.

Tse, P. (2022). Metaphysical idealism revisited. *Philosophy Compass* ***17***(7): 1–21.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind.* MIT Press.

Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press.

Uzgalis, W. (2001/2022). John Locke. *Stanford Encyclopedia of Philosophy*. In E. N. Zalta (ed.), URL = < https://plato.stanford.edu/entries/locke/>

de Vignemont, F. (2018). The Extended Body Hypothesis: Referred Sensations from Tools to Peripersonal Space. In A. Newen, L. De Bruin, and S. Gallagher (eds.), *The Oxford Handbook of 4E Cognition* (pp. 389–404).

Wagner, M. (1985). The Metric of Visual Space. *Perception & Psychophysics* ***38***(6): 483–495.

Walton, K. L. (1990). *Mimesis as make-believe: on the foundations of the representational arts*. Cambridge, Mass.: Harvard University Press.

Wildman, N. & McDonnell, N. (2020). The puzzle of virtual theft. *Analysis* ***80***(3):493–499.

Williams, J. R. G. (2020). *The Metaphysics of Representation.* Oxford: Oxford University Press.

Yablo, S. (1992). Mental causation. *Philosophical Review* ***101***(2): 245–280.

Yablo, S. (1995). Singling out properties. *Philosophical Perspectives* ***9***: 477–502.

Zeimbekis, J. & Raftopoulos, A. (2015). *The Cognitive Penetrability of Perception: New Philosophical Perspectives.* Oxford: Oxford University Press.