

MA Research Report:
Van Inwagen's Modal Skepticism¹.

Peter Hawke

Under the supervision of Dr. David Martens
Department of Philosophy, School of Social Sciences
University of the Witwatersrand
Private Bag 3, P O WITS 2050, South Africa

August 8, 2008

¹A research report submitted to the Faculty of Humanities, University of the Witwatersrand, Johannesburg, in fulfillment of the requirements for the degree of Master of Arts, by research report and coursework.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Arts in the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination in any other university.

Peter Michael Hawke

Date

Abstract

In this research report, the author defends Peter van Inwagen's modal skepticism. Van Inwagen accepts that we have much basic, everyday modal knowledge, but denies that we have the capacity to justify philosophically interesting modal claims that are far removed from this basic knowledge. The author also defends the argument by means of which van Inwagen supports his modal skepticism. Van Inwagen argues that Stephen Yablo's recent and influential account of the relationship between conceivability and possibility supports his skeptical claims. The author's defence involves a creative interpretation and development of Yablo's account, which results in a recursive account of modal epistemology, what the author calls the "safe explanation" model of modal epistemology. The defence of van Inwagen's argument also involves a rebuttal to objections offered to van Inwagen by Geirrrson and Sosa.

Contents

1	Introduction	1
2	Preliminaries	6
2.1	Qualified possibility	6
2.2	Conceivability and possibility	9
2.3	Conceivability-possibility arguments, strong and weak	10
2.4	Sorcery, ghosts, ghouls and goblins	14
2.5	The first type of modal skepticism	17
2.6	What is conceivability?	18
2.7	Imagination	19
2.8	Skeptical worries about the imagination	22
2.9	Possible worlds and fictional worlds	24
2.10	The second type of modal skepticism	27
3	Van Inwagen’s Skepticism	29
3.1	Yablo’s conceivability	31
3.2	Van Inwagen’s skeptical argument	35
3.3	Van Inwagen’s argument formalized	37
3.4	Objections to van Inwagen’s argument	39
3.4.1	Objection to the first premise	39
3.4.2	Objection to the second premise	43

3.4.3	Objection to the third premise	45
4	The “Safe Explanation” Model	46
4.1	An analysis of Yablo’s conceivability	47
4.1.1	What is it to imagine a world?	48
4.1.2	What is it to verify a proposition within an imagined world?	54
4.2	The “safe explanation” model of modal epistemology	58
4.2.1	Believability and the blueprint paradigm	60
4.2.2	Explanations, infinite regress and circularity	65
4.2.3	Defending the similarity principle	70
4.3	Conclusion	71
5	Skepticism Reconsidered	73
5.1	Can we acquire modal knowledge without Yablo-conceivability?	75
5.2	Is Yablo-conceivability too demanding?	77
5.3	Van Inwagen’s argument revised	80
5.4	Are we able to Yablo-conceive of ghosts, ghouls, goblins and sorcery?	81
5.4.1	Ghosts	82
5.4.2	Ghouls	84
5.4.3	Goblins	85
5.4.4	Sorcery	88
6	Conclusion	92
6.1	Summary of results	92
6.2	Avenues for further research	94

Chapter 1

Introduction

Modal questions are those that concern possibility (that is, what could have been) and necessity (what could not have been otherwise). A number of long-standing philosophical puzzles have arisen from the modal domain. Indeed, grappling with such modal issues has been a prominent feature of philosophical discourse throughout the history of Western philosophy. Aristotle and the medieval logicians were concerned with modal problems, as were prestigious philosophers of the modern period such as Descartes, Arnauld, Hume and Reid. In recent decades, interest in problems of modality has become particularly focused, thanks to the contributions of prominent philosophers such as Willard Quine, David Lewis, Alvin Plantinga and Saul Kripke to this area.

The myriad difficulties that arise in connection with modal claims can be introduced as follows. Consider the claim “it is possible that this copy of Locke’s *Essay* be open at page 231” (I’m referring here to my copy of the *Essay*, which, as I write, is lying on my desk). Now, as with all coherent statements, this claim is either true or false¹. It is true precisely when the situation it describes is in agreement with the facts and false precisely when this situation is not in agreement with the facts. Fur-

¹I will assume, for the duration of this paper, that statements of modality are either true or false, a position in line with common sense.

thermore, I am certain that we all agree that this modal statement is true (you'll have to take my word for it that no pages are missing from this copy, pages haven't been glued together, and so on). However, having agreed to this banal truth, we might find ourselves surprised to register difficulty in answering the following question: in virtue of *what facts* is this statement true? It is easy to see which facts would support the truth of "this copy of Locke's *Essay* is open at page 231" - this statement is true precisely when the book is *in fact* open at page 231 and false otherwise. What is more, if this forgoing non-modal statement is true, it seems reasonable enough to conclude the truth of our original modal statement, according to the indisputable principle that the actuality of a state of affairs implies its possibility. So, in this case, the fact that the book is open at page 231, a fact which can be uncontroversially verified, supports the modal claim. However, as it happens, it is, in actual fact, *false* that the book is open at page 231 (the book is, in actual fact, closed). Nevertheless, it remains true that the book *could* be open to that page. Suddenly, it has become mysterious what fact grounds this truth. What aspect of reality does the modal proposition accurately represent, given that the non-modal proposition embedded in our modal statement does not match up with the actual world? Are we dealing with a different notion of truth here? And how do we come by these truths, if they can't be verified through experience?

Hence, the most general question one might ask in the philosophical study of modality is this: in virtue of what is the statement "it is possible that p" true, where p is some proposition? This question can be understood in a metaphysical sense. On this front, it has become common-place for philosophers to suggest that asserting the possibility of p is equivalent to claiming the existence of a *possible world* of which p is true. The challenge then becomes to give a satisfying account of the metaphysical and ontological status of possible worlds. However, our central question can also be read in an *epistemological* light: when are we *justified* in asserting that p is possible?

This paper deals with modal epistemology. The core issue is thus one of justification. One might object to dealing with epistemic issues while the related metaphysical, ontological and semantic issues remain disputed and somewhat mysterious. How can we hope, one might say, to ascertain how we *know* things about possibility and necessity while it is still controversial what we are talking about in the first place? Surely, a successful resolution to the metaphysical issues will lay the groundwork for the theory of epistemology? In reply, I suggest that, as it stands, an investigation into the epistemology of possible worlds is just as likely to throw light on the metaphysical situation as the other way round. I will assume for the duration of this paper that we are competent language users who utilize the concepts of modality with relative success, and, furthermore, that we know the truth of many modal statements. These appear to me to constitute reasonable assumptions. Given these assumptions, any metaphysical theory will have to fit this epistemic starting point: if a metaphysical theory entails that we don't know the basic, banal modal statements that allow for daily living, then so much the worse for that metaphysical theory. Thus, it is legitimate to begin an investigation into modality leaving the metaphysical details hanging, while instead asking how we know what we do know about possibility and necessity, and what the extent of this knowledge is. Along the way, we might learn more about what a successful theory of possible worlds needs to involve.

The goal I have in mind for this paper, in studying modal epistemology, is fairly specific - namely, to explore the *scope* of our capacity to justify modal belief. When it comes to statements of the form "it is possible that p", we no doubt feel certain that some adequate answer to the question of justification can be found (whatever it is) if the proposition p is such that the state of affairs it describes is one that is familiar to everyday life. For instance, it seems entirely possible that I could have had a cup of tea this morning, as opposed to the glass of orange juice I did in fact drink. If anyone were to deny the truth of such a banal statement, or even refuse

to commit to its truth, we would probably label them as dysfunctional. However, what of the possibility-claims that *philosophers* are fond of making, those that are far removed from everyday life? Important and influential philosophical arguments, with profound conclusions, have hinged on premises such as “it is possible that a mind can exist disembodied”, “it is possible that pain can occur without c-fibers firing”² and “it is possible that a perfect being exists”. The truth of these propositions, and the conclusions they support, is controversial, to say the least. Philosophers have happily pushed assertions of possibility and necessity to their limits for centuries. This gung-ho approach invites skepticism: not skepticism aimed at everyday, or even scientific, modal claims, but rather at the maverick pronouncements of the philosopher. I intend to defend this skeptical outlook in this paper, founding it upon a reasonably detailed theory of the epistemology of possibility, one that hopefully shows some promise in being developed in the direction of a complete account of modal epistemology (that is, one that also deals with the epistemology of necessity and impossibility).

I will launch this investigation within a very specific context. Modal skepticism, of the ilk I have mentioned, has recently been discussed in the literature. In his 1998 paper “Modal Epistemology” [27], Peter van Inwagen expresses doubt that extravagant philosophical claims about possibility, such as those mentioned, can adequately be justified, although he is happy to regard our normal, everyday possibility-claims as justified. My interest in this paper stems chiefly from its final section, where van Inwagen argues that his brand of modal skepticism is supported by a recent and compelling account of the basis of our modal knowledge due to Stephen Yablo [27, pp.76-81]. Yablo argues in favour of the hypothesis that conceivability acts a guide to possibility, where he takes the conceivability of p to mean that one “can imagine a world that [one] take[s] to verify p ” [29, p.29]. Like van Inwagen, I feel Yablo’s theory has much to recommend it. Thus, the central question I intend to answer in

²Assuming, of course, that pain is in fact accompanied by c-fibers firing. I’m not certain what the latest scientific research on this is.

this paper is the following: does Yablo's account support modal skepticism as van Inwagen claims? The answer, I think, is "yes", but in order to see this, Yablo's account, and what van Inwagen makes of this account, will need to be clarified, and various objections will need to be dealt with. In the end, what I hope will emerge is an expansion of Yablo and van Inwagen's ideas, developed further in the direction of a full-blown theory of modal epistemology, one which finds a place for both our most basic modal beliefs and our more advanced modal beliefs, finds a role for both the imagination and rational principles in modal inquiry and rejects the lofty modal claims made by some of philosophy's most ambitious, interesting, but unfortunately sorely unpersuasive, arguments.

To begin, we need an account of Yablo's theory of conceivability and van Inwagen's ensuing argument for modal skepticism. This will require a more extensive background on modal issues, which I will now provide.

Chapter 2

Preliminaries

2.1 Qualified possibility

A vast number of different types of *qualified* possibility have cropped up in the philosophical literature in recent decades, including epistemic possibility, conceptual possibility, logical possibility, mathematical possibility and nomological possibility (which includes biological and physical possibility)¹. To avoid confusion, I make the following stipulation: when I use the word “possible” in this paper I am referring to possibility *simpliciter* or possibility *period*, what has come to be known, following Kripke [10], as *metaphysical possibility*. What is metaphysical possibility? This is the million-dollar question. Here are some alternative formulations, however:

p is metaphysically possible iff² it could have been the case that p iff there exists a possible world such that p is true of that world.

What then is meant by a *qualified* modality? A proposition p is possible, according to some qualification, just in case it is possible in conjunction with certain given

¹The typology I present here is generally accepted amongst philosophers, but aspects of it remain controversial. For a discussion of the view that metaphysical and physical possibility (necessity) coincide, see [19], [4] and [12]. For discussion concerning the relationship between logical and metaphysical possibility (necessity) see [18], [1], [2] and [6].

²As is the norm, I will abbreviate the logical connective “if, and only if” by “iff”.

restrictions. In other words, p is possible, in a qualified sense, if it is true of at least one possible world in a subset of all possible worlds, where this subset is defined by stipulating that those are the possible worlds in which certain conditions hold. To illustrate this, let us consider the two most commonly utilized (so-called) qualified modalities, they also being important to this paper³:

physical possibility- a proposition p is physically possible just in case it is possible in conjunction with the (actual) laws of physics. In other words, p is physically possible just in case it is true of at least one possible world where the laws of physics are true of that possible world. This type of possibility is obviously of particular interest to scientists and engineers. Other kinds of nomological possibility, such as chemical possibility and biological possibility may well be further restrictions on physical possibility. This is also the type of possibility that the layman is most likely to confuse with metaphysical possibility. However, many philosophers feel it is a mistake to conflate physical and metaphysical possibility, as, for one thing, they argue that the laws of nature could have been different. Thus, some philosophers hold that it is metaphysically possible that two bodies are not attracted to each other by a force proportional to the inverse of the square of the distance between the two bodies. On the other hand, this scenario may not be physically possible.

logical possibility - a proposition p is logically possible just in case it does not generate a contradiction, either explicitly (such as propositions of the form “ $q \& \neg q$ ”) or implicitly. A *set* of propositions is logically possible if the member propositions of that set are mutually consistent (according to some system of formal logic). This is an interesting category of possibility, since it isn’t really clear that it is a category of possibility *at all*, despite philosophers often speaking

³See Kripke [10] for Kripke’s famous discussion concerning epistemic possibility. See Yablo [30] for an interesting discussion concerning the notion of “conceptual” possibility.

this way. For one thing, logical “possibility” is sometimes claimed to be a *less* restricted class than metaphysical possibility. Since metaphysical possibility is possibility *simpliciter*, this can only amount to some logical possibilities being “possibilities” in name only⁴. For instance, “Peter Hawke is, and always has been, an inert pile of rock” appears to be metaphysically impossible, given that Peter Hawke’s status as a human being (at least at his origin) is a property he holds essentially (as Kripke [10], and many others, hold). Yet, there is nothing immediately contradictory about such a statement and thus nothing logically suspect. Of course, a contradiction *can* be generated if the statement is coupled with a list of Peter Hawke’s essential properties, which may be enough to qualify the statement as a logical impossibility. Formal logic, by itself, cannot generate such a list, however - this requires an insight into metaphysical possibility that goes beyond the mere manipulation of symbols. Or so it seems. Hence, some, including van Inwagen [27], have suggested that logical considerations can only act as a reliable guide to *impossibility* (it is generally accepted that contradictions can immediately be ruled as impossible⁵) and *necessity*, while consistency does not count as conclusive evidence for possibility.

⁴Of course, one way to avoid this outcome is to deny that metaphysical possibility (necessity) is possibility (necessity) *simpliciter*. That is, one can deny that metaphysical possibility is equivalent to unrestricted possibility. This position presents two problems. In the first place, there is difficulty in understanding what is supposed to be meant by a “metaphysical” restriction on possibility. Indeed, it seems to me that treating metaphysical possibility as a restricted possibility is likely to end up with metaphysical possibility dissolving into one of the other types of possibility (be it logical, conceptual or whatever). In the second place, treating metaphysical possibility as a restricted type of possibility will still leave us with the problem of making sense of possibility *simpliciter* (that is, the concept which metaphysical possibility is meant to be a restriction on). All in all, I think the correct spirit in which to view metaphysical possibility, and the accompanying talk of possible worlds, is as an attempt to understand possibility *simpliciter*. Hence, I will assume that metaphysical possibility is possibility *simpliciter* throughout this paper (indeed, the way I introduced metaphysical possibility, I have explicitly identified the two).

⁵Despite this general acceptance, it is worth noting that there are philosophers that think some contradictions *are* possible (even actual). This position is known as *dialetheism*. See [14].

2.2 Conceivability and possibility

What leads us to assert claims of possibility and necessity? To show a proposed state of affairs to be impossible or necessary, we often utilize pure reason: an inconsistency in a proposition demonstrates its impossibility; a proof for a mathematical statement demonstrates its necessity, and so on. How, though, do we normally conclude that a state of affairs is merely *possible*? One obvious answer, it seems (given that this is an answer often proposed by both philosopher and layman alike), is that we conclude the possibility of a state of affairs if we can clearly conceive of a scenario in which that state of affairs is instantiated. For instance, we conclude that Al Gore could have won the 2001 election in the United States because we can, seemingly without effort, conceive of a scenario in which he did win (perhaps in this scenario, Ralph Nader decides not to run and doesn't end up stealing votes from Gore in Florida). This technique appears to be at work even when we discern the most humdrum possibilities. How do I know that the furniture could be rearranged in my living room? Well, it certainly seems easy enough to *conceive* of the furniture being laid out in a variety of different patterns.

The seeming connection between conceivability and possibility has long been exploited by philosophers. David Hume, for instance, was convinced that conceivability *implied* possibility, as is captured in this well-known quotation:

'Tis an established maxim in metaphysics, that whatever the mind clearly conceives includes the idea of possible existence, or, in other words, that nothing we imagine is absolutely impossible [9].

Hume, famously, uses this principle to conclude that a cause is neither necessary nor sufficient for its effect. Since one can quite easily conceive of the effect occurring without the cause and the cause occurring without the effect (one can easily imagine a different effect following the cause, for instance), Hume concludes that no necessity

is involved with the cause and effect relationship. This is an argument that has been taken seriously by philosophers. Hume, were he alive, would no doubt be pleased to see that it has retained its persuasive force to this day.

An equally famous example is that of Descartes' argument for mind-body dualism in the *Meditations* [3]. Here, Descartes moves rapidly from the fact that he can clearly (and distinctly, as he puts it) conceive of his mind existing in a disembodied state, to the conclusion that his mind and body *could* exist separately (and, hence, that his mind and body are not identical).

More recently, philosophers and physicists have made ample use of thought experiments, whereby a conclusion is established through the description of an imagined scenario. Thought experiments are a standard tool in the analytic philosopher's arsenal: this technique has found application in every field of philosophy, to the extent that it is safe to say that analytic philosophy would be rendered unrecognizable without it. On the scientific side, scientists as prestigious as Galileo and Einstein have freely made use of thought experiments, often to great acclaim. Having noted this illustrious history, it must be acknowledged that thought experiments, seemingly, rely upon the principle that conceivability is evidence for possibility: in many cases (perhaps all), a thought experiment loses much or all of its force if it is admitted that the described scenario is, in fact, *impossible*⁶.

2.3 Conceivability-possibility arguments, strong and weak

Having established the seeming importance of conceivability in arriving at our modal judgments, let us now frame such arguments a little more formally.

⁶See Sorensen's *Thought Experiments* [23] for a thought-provoking survey and analysis of the use of thought experiments in rational inquiry.

The strong conceivability-possibility argument

- P1. It is conceivable that p.
- P2. If it is conceivable that p, then it is possible that p (**strong c-p principle**).
- C. So, it is possible that p.

I call this version of the argument the *strong* conceivability-possibility argument, since the conclusion of the argument is, straightforwardly, that p is possible. The second premise entails that it is not the case that one can conceive of p while p is impossible. A weaker version of this type of argument can be stated by weakening the second premise.

The weak conceivability-possibility argument

- P1. It is conceivable that p.
- P2. If it is conceivable that p, then one is justified in asserting that it is possible that p, in the absence of evidence to the contrary (**weak c-p principle**).
- C. So, one is justified in asserting that it is possible that p.

As we all know, being justified in asserting the truth of p does *not* entail the truth of p (although the two are, of course, compatible). On the other hand, if one is guaranteed the truth of p (that is, it is known that p), then one is obviously justified in asserting the truth of p. Hence, the conclusion of this second argument is weaker than the previous argument.

Let us for a moment consider the common first premise of these two arguments. The statement “it is conceivable that p” is ambiguous, so let us settle on the most suitable interpretation for our purposes. “It is conceivable that p” can be paraphrased as “p can be conceived of”. Now, this statement is ambiguous, in the first instance,

because it does not make clear *who* needs to be able to conceive of p in order for the statement to be true. Ought the statement be read as “anyone can conceive of p ”? Or ought it be read as “at least one person can conceive of p ”? How the statement is read in terms of these two options makes a difference to conceivability-possibility arguments, since the power to conceive is an ability, or capacity, that varies from person to person. In the first place, it appears that what a particular person can conceive of depends upon that person’s body of knowledge. For instance, it might well be the case that a physicist can conceive of certain physical scenarios that a person completely ignorant of physics could not conceive of, simply because the latter person does not have the theoretical background of the former. In the second place, it might be the case that the ability to conceive is a capacity that varies according to the inherent intellect of a person. It would not be surprising, I think, to discover that a person regarded as having a genius-level intellect is able to conceive of certain scenarios that a person of below-average intellect fails to conceive of (even after a great deal of coaxing and explanation by the genius). There is also a second ambiguity that arises from the statement “ p can be conceived of”: ought this statement be read as “ p *has* been conceived of” (or, even, “ p is being conceived of”) or ought it be read as “ p *could* be conceived of” (that is, “it is possible that p be conceived of”)? The first interpretation can be taken to state that, as a matter of actual fact, someone has conceived of p at some point in the past. The second interpretation can be taken to state that it is a possibility that someone conceives of p . This second statement can be true even if no-one in actual fact *has* conceived of p .

How then to interpret the statement “ p can be conceived of”? First, it is clearly far too demanding to insist, given the variation in our capacity to conceive, that everyone can conceive of p . If it is known that at least one person can conceive of p , this ought to suffice for the purposes of supporting a conceivability-possibility conclusion. With

regard to the second ambiguity, it is, for the purposes of conceivability-possibility arguments, advisable to phrase the antecedent of the c-p principle in a manner independent of the concept of “possibility”, in order to safely avoid a vicious regress (it may be noted, of course, that “p has been conceived of” implies “it is possible that p be conceived of”). In total, “it is conceivable that p”, as it appears in the above two arguments, can be spelt out as “at least one person has, in actual fact, conceived of p”.

Let us now take a closer look at the weak c-p principle. The strong and weak c-p principles are obviously compatible. However, there is no inconsistency in the strong c-p principle being false and yet the weak c-p principle being true. This is generally the spirit in which the weak c-p principle is defended: as the basis for saving conceivability-possibility arguments in the face of serious objections to the strong c-p principle. On a similar note, the weak c-p principle opens up the possible line of argument that while conceivability acts as evidence in favour of possibility, conceivability does not count as *conclusive* evidence. The weak c-p principle remains true if conceivability is merely a reliable guide to possibility, albeit not an infallible guide. This in turn opens the way for the idea that conceivability can act as better or worse evidence, depending on what is conceived and how it is conceived. In other words, the potential for *degrees* of evidence finds a natural place in the thought of those who defend the weak c-p principle, yet reject the strong c-p principle.

Both the strong and the weak c-p principles present immediate respective disadvantages. The strong c-p principle makes a bold claim, and thus opens itself up for counter-examples. The weak c-p principle, on the other hand, if defended in conjunction with a rejection of the strong c-p principle, leads to a demand for an account of modal error: if conceivability is a *fallible* guide to possibility, how and why can it fail?

In this paper, I will chiefly be interested in the weak conceivability-possibility argument, as this is the version that Yablo's account supports and, I believe, the best hope for defending a conceivability-possibility argument.

2.4 Sorcery, ghosts, ghouls and goblins

I think it is now sufficiently established that conceivability-possibility arguments are advertised as playing an important role in daily thought and rational inquiry. Philosophers have not hesitated to push this technique to its limits, however, producing possibility-conclusions that are startling both for their sweep and their counter-intuitiveness. Indeed, it would be surprising to meet someone who believes some such possibility conclusions, who is neither familiar with philosophical thinking, nor a child below the age of ten.

I will now present four conceivability-possibility arguments that will be of importance to this paper. Each argument presents a possibility conclusion that (in one form or another) is a) important to philosophy and accepted by some esteemed philosophers and b) I think, controversial. One of the central goals of this paper is, in time, to assess these arguments.

I will phrase these possibility-arguments in terms of ghosts, ghouls, goblins and sorcery, the stuff of fantasy novels and fairy tales, in order to emphasize the element of both the dramatic and the fantastic displayed by these conclusions. Furthermore, I think there is little *prima facie* doubt that these things can be conceived of and conceived of in detail: they have been the subject of lively and vivid fiction for many centuries!

A few preliminaries: a *ghost* is an immaterial, non-corporeal, conscious being, generally depicted as a soul separated from its body. A *ghoul* I will take to be equivalent to a zombie: a being whose body operates, but has no mind and thus no

conscious states (in fiction, they tend to distinguish ghouls from persons by way of their - murderous - behaviour, but this is hardly necessary: the ghoul's behaviour may match an ordinary person's). A *goblin* is similar to a human being, but with a few key differences: a goblin is small, green-skinned and pointy-eared. Goblins are also, as a rule, somewhat cunning and have a vicious demeanor. Finally, a world in which the potential for *sorcery* exists, I stipulate, is one in which the following law (the "alakazam" law/spell) holds: if one holds some wormwood in one's left hand, points at an object with one's right index finger and utters the word "alakazam!", then whatever is in the line of sight of one's finger will promptly burst into flames. Here, then, are the arguments (I have phrased them in the form of strong conceivability-possibility arguments. They can, of course, be easily rewritten as weak conceivability-possibility arguments):

Argument 1: The Ghost Argument

- P1. It is easy to clearly and distinctly conceive of a situation in which a ghost exists.
- P2. If something can be clearly and distinctly conceived, then it is possible.
- C. So, it is possible that ghosts exist.

Argument 2: The Ghoul Argument

- P1. It is easy to clearly and distinctly conceive of a situation in which a ghoul exists.
- P2. If something can be clearly and distinctly conceived, then it is possible.
- C. So, it is possible that ghouls exist.

Argument 3: The Goblin Argument

- P1. It is easy to clearly and distinctly conceive of a situation in which a goblin exists.

P2. If something can be clearly and distinctly conceived, then it is possible.

C. So, it is possible that goblins exist.

Argument 4: The Sorcery Argument

P1. It is easy to clearly and distinctly conceive of a situation in which the alakazam law holds.

P2. If something can be clearly and distinctly conceived, then it is possible.

C. So, it is possible that the alakazam law holds.

Despite the fantastical quality of these conclusions, some philosophers have taken them very seriously. Argument 1 is crucial to the argument for cartesian dualism. Argument 2 forms the basis for a more recent argument in favour of dualism [25]: if it is possible that our bodies operate precisely as they do and yet no minds are “housed” in our bodies, then it appears that mind and body cannot be identified. As for Argument 3: philosophers are fond of asserting the possible existence of beings such as unicorns, naturally purple sheep, utility monsters and God. The nature of such beings is normally stipulated by the philosopher according to a (partial) list of essential qualities that that being has (a unicorn is like a horse, but with a naturally occurring horn protruding from its forehead; God is a being that is omnipotent, omniscient and omni-benevolent; and so on). It is generally allowed that such creatures are conceivable. The apparent conceivability of God has, in particular, been exploited by philosophers, by way of the well-known family of ontological arguments in the philosophy of religion. For my representative argument of this type, I have chosen, as opposed to asserting the possibility of God, the more modest conclusion that goblins might have existed. Finally, Argument 4 is merely an extension of Hume’s conclusions concerning the contingency of cause and effect, and of the laws of nature. If different

laws of nature are possible, surely the alakazam law is a suitable candidate for a possible law of nature?

While some philosophers have accepted such conclusions, others have been more suspicious. Indeed, the above arguments pose the following challenge: if one is loathe to accept the conclusion of one of these arguments, where does the flaw in the argument lie? Is one forced to abandon the c-p principle, weak and strong?

2.5 The first type of modal skepticism

Our reliance upon conceivability to make modal conclusions, and the appearance of controversial possibility-arguments on the scene, has provoked at least two types of modal skepticism from philosophers. I will discuss the first type here. Some philosophers deny the second premise of the conceivability-possibility argument. That is, some philosophers deny that conceivability can act as a guide to modal possibility *at all*. This, quite clearly, places a great deal of our modal knowledge in jeopardy (even if the skeptic offers some alternative to conceivability). Such philosophers include Thomas Reid [17], John Stuart Mill [13] and, more recently, Paul Tidman [26]. Mill expresses a general worry when he says that

our capacity or incapacity of conceiving a thing has very little to do with the possibility of the thing in itself; but is in truth very much an affair of accident, and depends upon the past history and habits of our own minds [13].

In short, the concern is that the domains of conceivability and possibility appear to be entirely separate. What counts as conceivable is exceedingly subjective: it has everything to do with the state of one's mind. Thus, conceivability is constrained in important ways which seem to have nothing to do with what is possible. In other ways, however, conceivability is considered to be worryingly broad in scope: human

beings can conceive of all kinds of crazy things. On the other hand, what counts as possible is an entirely objective matter concerning the nature of reality. What connection could the two possibly share?

Recently fueling this debate has been the work of Saul Kripke[10] and Hilary Putnam[15], both of whom suggest that the existence of *a posteriori* necessary truths makes it all too easy for us to conceive of situations that are impossible. Putnam, famously, suggests that we have no trouble conceiving of a situation in which we discover that water isn't H_2O (merely imagine finding evidence that water is composed of XYZ, as opposed to H_2O), although it has turned out, following an extensive scientific investigation, that water is essentially H_2O . This leads Putnam to note

Conceivability is no proof of logical possibility [15].

This type of skepticism is not really my concern in this paper. I will, however, offer some reason to safely set it aside, given both the obvious seductiveness of this position and the obvious futility of pursuing my central project in this paper if this brand of skepticism is correct. In order to do this, however, we will have to raise an important issue that I have thus far avoided. What exactly is meant by “conceivability”?

2.6 What is conceivability?

When it is suggested that conceivability acts as a guide to possibility, what sense of “conceivable” are we making reference to? If one finds a proposition p conceivable, does this mean that one finds it understandable? That it describes a state of affairs that cannot be ruled out as actual, on the basis of the given evidence? That one can imagine a situation in which one is convinced of the truth of the proposition? That one can imagine a situation in which the proposition *is* true? That one can visualize the proposition?

The word “conceivable” is surprisingly fluid, and one needs to be wary of its meaning in conversation and philosophical discourse. For my part, it seems clear that the conceivability that is relevant to possibility arguments is *imaginability*. This, it seems to me, is what most people mean when they claim possibility on the basis of conceivability: to conceive of a proposition, in the sense relevant to conceivability-possibility arguments, is to imagine a situation of which that proposition is an unequivocally accurate description (this, in fact, is close to Yablo’s suggestion). Furthermore, I think any other interpretation of conceivability will quickly run into trouble as a guide to possibility, although I won’t defend this claim here.

2.7 Imagination

Let us thus agree that our topic is the imagination. What, however, is the imagination? One might be tempted to identify “imagining that p” with “visualizing that p”. This identification is over-easy and a little rash, however, despite it being clear that visualization often plays a role in imagining. There are clear cases of imaginings that are not visualizations. For instance, when I sit down to finally write my crime noir novel, I might imagine that it is fictionally true that the lead character Nick Nail used to have a drinking problem (in the fictional past - when the book’s narrative begins, he has cleaned up). I might also imagine that Nick thinks to himself “this dame is trouble”. Neither of these imaginings, it seems to me, involves any visualization. I don’t visualize that Nick had the drinking problem (I don’t bother to visualize his past) - nevertheless, I imagine him with this property - and I certainly don’t visualize his thoughts (although I might visualize him sitting with a thoughtful look on his face). Indeed, it turns out to be a difficult and important question to say precisely what the imagination is, one that I will not attempt to answer in this paper (beyond some suggestive comments). I will instead proceed by making a few (I hope

reasonably harmless) assumptions about the nature of the imagination, and the role that it plays in our lives. I trust the reader will agree that all of these assumptions are *prima facie* plausible and in line with common sense, even if some aspects of them are somewhat vague and even if all, I am sure, require a proper philosophical grounding. In total, my assumptions about the imagination are meant only to assert a close link between the imagination and fictional worlds (or, at least, those fictional worlds which are worth talking about).

First, I assume that there are things called *fictional worlds*, with examples being the fictional worlds described in books such as Mary Shelley's *Frankenstein* and Aldous Huxley's *Brave New World*. I will not attempt to say precisely what a fictional world is (or whether this notion genuinely deserves the "world" moniker, given, as I think will become clear, that a fictional world is a very different kind of thing to a possible world, including the actual world). However, I will, at one point or another, outline some of the features of fictional worlds. In general though, I trust that the reader has a working understanding of the notion of a fictional world, with which I can proceed. Second, I assume that fictional worlds are constructions. That is, I assume that they are *made* (constructed) by persons. This construction involves, in one way or another, setting and representing the features of the fictional world. This construction can be purely *mental* in nature, as when one visualizes a scene in one's mind. It can also involve physical acts, as when one sets and represents the features of a fictional world in a painting. Fictional worlds can thus be shared when expressed in films, books or other media⁷. A person responsible for the construction of a fictional world is called an *author* of that fictional world. Every fictional world has at least one author, who is responsible for its construction. Third, I assume that the imagination can be thought of as a mental capacity or ability or "power", that occurs

⁷I will not discuss the interesting question as to whether media such as literature genuinely *share* already existing fictional worlds, or if they rather *prompt* an audience member to construct their *own* fictional world, albeit with certain prescribed features.

in various degrees across different people. By a “capacity”, I mean that the imagination allows a person to perform certain kinds of mental acts (“imaginative” acts). Finally, I assume that this capacity is necessary for a person to construct interesting fictional worlds (that is, constructing an interesting fictional world is an imaginative act). Indeed, I will assume that if someone lacked imagination completely, then they would be unable to construct interesting fictional worlds. I do not, however, claim that imagination is sufficient for constructing an interesting fictional world. I make no attempt to rigorously clarify the notion of an “interesting” fictional world. Rather, I appeal to the following intuitive distinction: on the one hand, I think it is generally agreed that the world depicted in a great work of literature, such as *Frankenstein*, is “worthwhile” or “interesting”. On the other hand, if I were to construct a fictional world purely by way of some mechanical procedure (for instance, I form the fictional facts of that world by picking sentence parts - subjects, predicates and so on - out of a hat, and randomly putting them together to form a list of fictional facts), if indeed this counts as a fictional world, then I think it can hardly be expected for this fictional world to count as worthwhile, or of interest. “Interest value”, of course, occurs in degrees: the world of *Frankenstein* is of a great deal of interest; the worlds of Barbara Cartland novels are much less of interest; a mechanically produced fictional world, as above, is predictably flagrantly uninteresting. Notice how, in general, we tend to correlate the level of interest attached to a fictional world with a perceived level of imaginative capacity on the part of the author, that has been exercised in the making of the fictional world.

Hence, I assume that it makes sense to think of imagination as a capacity that can be utilized by a person, sometimes with skill, to the end of producing certain kinds of products. This is not a strange way to talk about the imagination. We praise someone for being born with the gift of a great imagination, or praise someone - perhaps an artist- for having developed their imagination to its “full potential”. We

say someone has a “great deal of imagination” if, for instance, they produce detailed and captivating fictional worlds and, likewise, might say, in a derogatory fashion, that someone has “no imagination” if they are unable to construct even basic fictional worlds when prompted.

As a final comment about the imagination (for now), it is worth noting that there appear to be different levels of (conscious) control the author can exert on the creation of a fictional world. Some authors, such as the writers of fictional literature, can self-consciously dictate every specified aspect of their fictional world. On the other hand, the imagination has the capacity to operate in a non-deliberate manner too, as when we dream, allow our minds to wander or find our imagination producing content in response to sensory events (good theater, for instance, is said to stimulate the imagination)⁸. In such cases, what the imagination produces can even *surprise* the author. There is also some kind of middle-ground: often one “commands” the imagination to produce a representation of some specific description, as when one attempts to spontaneously visualize a pirate walking across the deck of his ship, and finds oneself surprised at the precise features of the image that is produced.

2.8 Skeptical worries about the imagination

Having settled on the imagination as our topic, we can now sharpen the first modal skeptic’s objections to the c-p principle. In the first place, the modal skeptic can wheel out some tried and tested counterexamples against an imagination thesis. It appears that we can imagine fictional scenes of which it is unequivocally true that those scenes depict an impossibility. A favourite example is that of Escher’s “impossibility” pictures - his ever-ascending (ever-descending) staircase, for instance [5]. A

⁸See Walton [28] for an interesting discussion concerning “deliberate” versus “non-deliberate” imagination (indeed, these are Walton’s terms), amongst other insightful comments about the imagination. I do not claim, however, that the general views Walton expresses about the imagination match my own, nor are my views intended to match Walton’s.

slightly different kind of counter-example comes from the television show *Star Trek*: on the show, Captain Kirk and crew are regularly depicted as teleporting from one place to another. Now, not only is it in question whether it is physically possible for this to occur, it is also a philosophical problem whether one maintains one's self-identity if one's body is vaporized and reconstituted elsewhere. Thus, it is in question whether a person really *could* be teleported from one place to another. It is, however, an indisputable (fictional) fact of the series that Kirk does this all the time (it would seem absurd to argue that Kirk's adventures in the fictional world ends, despite appearances, the first time he steps into a teleporter. *Star Trek* chronicles *Kirk's* adventures, after all, not a copy of Kirk's. The show's authors, the real authority in such a world, quite clearly mean that *Kirk* arrives at his destination following teleportation, not a copy). Now, the point the skeptic can make is the following: surely being able to imagine the teleportation is not *in itself* enough to settle the question of possibility here. Something more is needed, which imagination does not provide. Finally, the skeptic can wheel out Kripke and Putnam's *a posteriori* impossibilities. I can surely imagine George W. Bush turning out to be an automaton - yet, assuming he is in fact human, this is a metaphysical impossibility.

These counterexamples will not concern me overly at this point. If they are fatal, they are fatal only to the strong c-p principle. The weak c-p principle, which is our real topic of discussion in this paper, does have *some* doubt cast upon it by such examples (they may demonstrate an unacceptable level of unreliability of the imagination), but they cannot, without further argument, bury the principle, as it is designed for the eventuality that conceivability might sometimes give the wrong answer. These counter-examples do, however, demand an explanation from the weak c-p supporter: does the imagination really fail us in these cases? And, if so, why?

A second avenue of attack by the skeptic presents a more pressing worry for the weak c-p principle. The skeptic is concerned that there is no reason to think there

is *any* connection between conceivability and possibility, whether the connection is fallible or not. As Paul Tidman states it:

The central question that must be answered to defend even this more modest proposal for conceivability [the weak c-p principle] is why should it be thought that there is *any connection at all* between what we can or cannot conceive and what is possible. These seem to be two distinct subject matters . . . Why should we even begin to think that anything that is possible is such that we can conceive of it? [26, p.306]

What answer can be given to this challenge? For one, I hope that the account of conceivability/imaginability that I defend in this paper will at least be developed enough and plausible enough to add some weight to the weak c-p principle. For now, however, I will only offer the obvious (and, some might say, too quick) reply to this worry: the use of imagination appears crucial to modal inquiry and there do not appear to be any alternative methods of investigation that could fill the gap if imagination-based methods are ruled out as irrelevant. Indeed, it appears to be the case that modal inquiry will be left unrecognizable without the inclusion of the imagination. This is no coincidence. The imagination appears to be *uniquely suited* to playing a role in granting us modal knowledge, whether or not it provides us with modal knowledge by itself, or in harmony with other factors. I will thus defend the claim that imagination is crucial (that is, a necessary condition) for useful modal inquiry in the next section, while laying some groundwork for the imagination-centered account of modal epistemology which will be developed throughout the paper.

2.9 Possible worlds and fictional worlds

If a possibility has been actualized, then our normal scientific methods of investigating the world can be utilized in discovering that possibility. On the other hand, the

difficulty in knowing non-actual possibilities, is that one is forced to depart from experience to do so. While experience of the actual world may *inform* or *suggest* certain non-actual possibilities, ultimately we must rely upon our *a priori* faculties to discern what is non-actually possible. Which *a priori* methods are suitable, however?

When investigating non-actual possibility, we do not generally ponder a single non-actual proposition, but rather a non-actual situation (in which a specific proposition may be embedded). Our goal is then to discern if this non-actual situation is in fact possible. How is it that we come to think of non-actual situations in the first place, however? The answer is obvious: we *make up* these situations. In other words, what we are scrutinizing in such an investigation is, in fact, a *fictional world*. When one is asked by a philosopher, for instance, to consider the possibility of p, or to consider a world in which certain states of affairs obtain (in the interests of a thought experiment, perhaps), what is expected of one is to use one's imagination to fill out a reasonably detailed imagined scenario in accordance with the instructions given by the philosopher. To this extent, it is *fictional* worlds that form the content of our musings about what is possible. On the face of it, this situation is not inappropriate. There is a certain practical necessity involved with utilizing fictional worlds in this regard: first, studying the actual world is of limited use when considering non-actual possibilities; second, it is not clear how we are meant to access non-actual possible worlds directly (on the other hand, we have unprecedented access to fictional worlds); third, methods that present an alternative (of equal power and flexibility) to constructing fictional worlds, for considering non-actual scenarios, are not obvious.

In total, it seems to me that fictional worlds, and by extension the imagination, is an obvious and promising tool for exploring possibility questions. Indeed, a central problem in the epistemology of possibility may be the following question: by what criteria can we safely conclude that a fictional world matches a possible world? In particular, are there some principles that need to be adhered to in the construction of

the fictional world that makes it more likely to coincide with a possible world? It seems to me that the answer to this last question is “yes”. One generally accepted principle is the *consistency* principle. Let us say that a fictional world is an *approximation* or *partial description* of a possible world if the set of fictional truths that completely describe the fictional world are a subset of the set of propositions that completely describe the possible world. If a fictional world is constructed in such a manner as to take care that no inconsistencies exist amongst its fictional truths, then it seems clear that this acts as *positive*, if inconclusive, evidence that a possible world exists for which that fictional world is a partial description. Are there other such principles that can act as further guides in the construction of fictional worlds? If there are, then, as one would expect, it will turn out that imagination *is* a guide to possibility - we simply have to make sure we imagine things *in the right way*. Failure to do so will result in the potential for modal error.

All this suggests an argument asserting the dependence of useful modal inquiry on the imagination:

- P1. Useful modal inquiry amongst humans occurs only if human beings can ponder non-actual scenarios.
- P2. For any person, if that person is pondering a non-actual scenario, then that person is pondering a fictional world.
- P3. If fictional worlds (or at least the kind of fictional worlds worth considering) have been constructed and pondered by human beings, then human beings have imagination.
- C. Therefore, if human beings did not have imagination, then useful modal inquiry would (at least largely) halt.

Let me end this discussion of fictional worlds and possible worlds with a quick direct comparison between fictional worlds and possible worlds. Possible worlds are

representative of the nature of reality; fictional worlds are representative of the imagination and the potential for creativity held therein. Possible worlds are discovered; fictional worlds are constructed. Possible worlds are consistent, insofar as the list of propositions that comprehensively describes a possible world will be mutually consistent; fictional worlds, as constructions, need not be consistent (we are all familiar with gaps in continuity in story-telling, or “holes in the plot”). Possible worlds are “complete”, in the sense that any proposition has a true or false value relative to a possible world; for any given fictional world, there are no doubt many propositions which are neither true nor false of that world, since the author has failed to construct the fictional world so that that proposition is fictionally true or false (if a character’s shoe size has not been set by the author, then it is neither true nor false that his shoe size is eleven). Finally, and importantly, our access to non-actual possible worlds is somewhat problematic and must be indirect; fictional worlds are routinely constructed, shared and enjoyed by human beings.

2.10 The second type of modal skepticism

At this point, I think enough has been said about the general worries of the first modal skeptic. There is, however, a second type of modal skepticism. This type will be the central topic of this essay. The second type of skepticism involves the *scope* of our powers to represent possibilities to ourselves. That is, this type of skepticism accepts that conceivability acts as a guide to possibility, but questions whether we can adequately conceive of scenarios that would support a great body of our possibility-claims. In other words, this skeptic accepts the c-p principle, and instead denies the truth of the first premise of the conceivability-possibility argument for some significant class of our modal beliefs.

Van Inwagen falls into this camp: he accepts that we have much basic, everyday

modal knowledge, but denies that we have the capacity to justify philosophically interesting modal claims that are far removed from this basic knowledge. Exploring this view will be the task of the rest of this essay.

Chapter 3

Van Inwagen's Skepticism

Van Inwagen is a skeptic of the second type. He appears happy to accept some version of the c-p principle. He is less enthusiastic to endorse far-out philosophical claims of possibility. Philosophers are prone to offer such claims with little or no supporting evidence: simply because they “see” the truth of such cases, or because of sloppily executed imaginative exercises. Van Inwagen sees such misplaced confidence as nothing more than a product of a philosophical culture that has grown accustomed to accepting such claims without question. The situation, van Inwagen suggests, is analogous to that when we make judgments of distance according to nothing more than the naked eye [27, pp.69-70]. When dealing with distances that are familiar to everyday living, normal eyesight and intuition are enough to act as reliable guides to distance judgments. One does not need special equipment or techniques to estimate that a room is four meters in diameter, for instance. However, when dealing with distances that are beyond our normal realm of experience, our intuitions can go horribly wrong (witness the tendency for ancient civilizations to wildly underestimate the distance to celestial bodies, even the sun or moon). Similarly, there is a class of everyday modal statements, what van Inwagen calls *basic modal statements*, that our modal intuitions (that is, the modal judgments we are prepared to strongly support in the absence of

clearly outlined reasons), honed by experience, can claim a certain authority over (or so van Inwagen claims). However, for more “remote” modal statements, those far removed from actuality, our intuitions cease to be of much use and more rigorous reasoning is required¹.

Van Inwagen’s position, as put forward in “Modal Epistemology”, is best introduced with a sample of both the modal claims which he accepts and those he rejects. Van Inwagen claims that we can know² the following possibility propositions, all of which he classifies as “basic”, everyday modal propositions:

- It is possible that the table be two feet to the left from where it in fact was.
- It is possible that John F. Kennedy should have died of natural causes.
- It is possible that the legs and top of this table need never have been joined together.

On the other hand, Van Inwagen is skeptical that we know the following possibility propositions:

- It is possible for there to be a perfect being.
- It is possible that I exist and nothing material exist.
- It is possible that there exist vast amounts of suffering for which there is no explanation.
- It is possible that there exist a naturally purple cow.
- It is possible that there exist transparent iron.

¹I will follow van Inwagen in using this vague terminology of “basic” and “non-basic” or “far-out” modal statements for the rest of the paper. The distinction being made here is intuitive, so it will be safe enough to apply it for the time being. Ultimately, however, I will attempt a more precise definition of “basic” and “non-basic”, as my account develops.

²Van Inwagen does indeed think we *know* these propositions [27]. Van Inwagen concedes, however, that he is not certain of the ground of this basic modal knowledge [27, p.74]. If he were being more cautious, as I will be, he might rather say that we justified in asserting the truth of such propositions.

The first three of the above propositions are, respectively, crucial premises in the argument for the ontological argument, cartesian dualism and the argument from evil. To support his claims, Van Inwagen argues that Stephen Yablo's recent and influential account of the relationship between conceivability and possibility supports his conclusions [27, pp.76-81].

3.1 Yablo's conceivability

Yablo argues that we have no reason to doubt conceivability as a generally reliable guide to possibility (although, cautiously, he admits that conceivability leaves room for some modal error) [29]. He achieves this by carefully separating out various senses of "conceivable" that are used in philosophical and everyday discourse. He concludes that only one sense of "conceivable" is a suitable candidate for what we mean when we say that conceivability guides possibility and, what is more, many of the objections that have been raised against a conceivability-possibility relation (including the Kripke-Putnam objection) fall away when we spell out what conceivability, in this case, amounts to. Here is Yablo's thesis:

p is [philosophically] conceivable for me iff[f] I can³ imagine a world that I take to verify p [29, p.29].

The alternative takes on "conceivability" that Yablo considers are the following:

- p is *conceivable_b* iff it is (not un)believable that p.
- p is *conceivable_{bp}* iff it is (not un)believable that *possibly*, p.

³One immediate complaint about Yablo's thesis is the modal ambiguity he introduces into his statement of the thesis through the use of words such as "can". Does "I can imagine a world" equate to "it is possible that I imagine a world"? In this case, there appears to be some threat of vicious regress, as this might suggest we require Yablo-conceivability to determine whether we can Yablo-conceive of p, which in turn will require Yablo-conceivability, and so on. To avoid this matter, I will attempt to remove this modally ambiguous language from Yablo's thesis when transplanting it to other contexts, such as when I formalize van Inwagen's argument.

- p is *conceivable_{ijb}* iff one can imagine justifiably believing that p .
- p is *conceivable_{itb}* iff one can imagine believing p truly.
- p is *conceivable_{ep}* iff one can imagine believing something true with one's actual p -thought [29, p.26].

What is notable about this list is, first, the subtlety in difference between some of these definitions and, second, the undeniable ease with which one can slip between these various senses of “conceivable” in everyday and philosophical discourse. *Conceivability_b*, which can be rephrased as “one finds no reason to rule out p as actual given the available evidence”, is particularly prevalent in laymen conversation.

Yablo holds that his “philosophical” conceivability is the only suitable candidate for supporting conceivability-possibility arguments, since it is the only candidate under review that passes the so-called “modal appearance test” [29, pp.4-6]. What is this test? Yablo argues that we need only concern ourselves with types of conceivability that *advertise* themselves as guides of possibility - that is, that provide the *appearance* that p is possible. For a type of conceivability to involve the appearance of possibility is for the state this conceivability produces to be one such that (i) this state is veridical only if possibly p (that is, my conceiving that p is true only if it is possible that p), and (ii) being in this state moves one to believe that possibly p . The manner in which conceivability ought to involve the appearance of possibility, is analogous to the manner in which perception involves the appearance of actuality, according to Yablo.

It is clear, upon inspection, that none of Yablo's alternative conceivabilities can be said to fulfill this criteria. Only his “philosophical” conceivability has any real hope of making the grade: the stipulation that the imagined world appear to “verify” p is tantamount to ensuring that the state be veridical only if a possible world exists of which p is true. Yablo thus sees as a consequence of this verification, that an

imagined scenario is only a verification of p if that scenario is not compatible with the falsity of p . In other words, the veridicality of the imagined scene must be such that it can only be fulfilled by p being the case. Yablo gives the example of imagining a super-computer print out a number n which is hailed as a counter-example to the Goldbach conjecture [29, p.31]. This imagining does not suffice as an example of Yablo's philosophical conceivability, and hence acts as no evidence in favour of the possibility of the Goldbach conjecture being false, since this scenario is compatible with the falsehood of the claim that there exists a counter-example to the Goldbach conjecture (the scenario can be filled out in the following way: no counter-example exists; the super-computer is flawed and expert observers have simply been too lazy to check its result).

Yablo furthers his case for his "philosophical" conceivability by arguing that key arguments against the weak c - p principle discreetly turn on "alternative" understandings of conceivability, losing their force when phrased in terms of philosophical conceivability. First, consider the following "circularity" argument: since unappreciated impossibilities are easy to conceive of, we always require *prior* modal knowledge in order to judge whether a conceived scenario is possible or not. Hence, conceivability is not the true source of our modal knowledge, as, on its own, it can only yield inconclusive data. Yablo, however, questions why a proponent of this argument is likely to hold that unappreciated impossibilities are so easy to conceive of (to the extent that this renders conceivability an unreliable tool). He decides that such a belief betrays a certain interpretation of conceivability - namely, conceivability_{bp}, or, "a conceivable proposition is just one not known to be impossible" [29, p.20]. On the other hand, the statistical hypothesis that "unappreciated impossibilities are almost always conceivable" is harder to swallow with Yablo's philosophical conceivability in mind. Indeed, in this case, there appears to be a strong correlation between impossibility and inconceivability. It seems we cannot (directly) imagine explicit contradictions or (directly)

imagine a situation in which two and two add up to three. Indeed, these kinds of situations are generally offered as typical examples of the unimaginable.

Yablo also considers the Kripke-Putnam *a posteriori* objection, from which he gleans an argument directed at the weak c-p principle [29, p.21], of similar bent to the previous objection, but restricted to the conceivability of *a posteriori* falsehoods.

P1. Whenever p is *a posteriori* false, I find it conceivable whether it is possible or not.

P2. Often, *a posteriori* falsehoods are impossible.

C. So, *a posteriori* falsehoods are often found conceivable despite their impossibility.

P2 is backed up by Kripke and Putnam's assertion that there are necessary *a posteriori* truths: that water is H₂O, that a person originated from particular genetic material, that Hesperus is identical to Phosphorus, and so on. Indeed, such truths appear to be numerous, and may, the objector suggests, be even more numerous than we think. P1 is backed up by the ease with which we can imagine finding out that the contrary holds to any *a posteriori* truth: just as I can imagine that the morning paper wasn't delivered this morning, so I can imagine that it was discovered that water is XYZ, not H₂O. The point of the objection is thus that *a posteriori* impossibilities are both numerous and easy to conceive and, thus, that conceivability-possibility arguments fail for the broad class of *a posteriori* falsehoods.

Yablo, however, questions whether the objector really has philosophical conceivability in mind when asserting P1. Indeed, Yablo thinks that Putnam and Kripke have either conceivability_{ijb} in mind here (one can imagine gaining evidence in favour of p), or conceivability_{ep} in mind (one can imagine a world in which one justifiably and truly believes a statement such as "water is not H₂O", but only because the meaning of this statement is different in that world - water there means "XYZ", not

“H₂O”). On the other hand, Yablo contends that we are *not* able to philosophically conceive of *a posteriori* impossibilities. For instance, one cannot imagine Hesperus and Phosphorus being distinct, because this would involve imagining Venus being distinct from Venus.

There is one other attractive feature of Yablo's account. Alongside his account of conceivability, Yablo offers the following account of inconceivability:

p is inconceivable for one iff one cannot imagine any world that one doesn't take to falsify p [29, p.29].

The advantage to this account is that inconceivability is not the denial of conceivability, opening the space for a third category: undecidability [29, p.31].

The possibility of p is undecidable for one iff one cannot conceive of p and yet p is not inconceivable to one.

3.2 Van Inwagen's skeptical argument

Assuming the (rough) truth of the above, van Inwagen argues in favour of the thesis that far-out philosophical possibility-claims are not Yablo-conceivable. Indeed, he thinks that the real status of such claims is *undecidable*, since he isn't happy to commit to Yablo-inconceivability either (although he doesn't argue for this last conclusion). Here, as far as I can see, is the argument van Inwagen puts forward, in rough form: in order to conceive of a p-world, one needs to imagine a world that one takes to *verify* p. However, this involves imagining a world in sufficient detail so that the scenario described (the one purported to constitute an imagining of the truth of p) cannot be compatible with the falsity of p. Now, where p is a philosophical claim⁴ far removed from the basic, uncontroversial modal claims we make as a matter of course

⁴By a “philosophical claim”, I mean, roughly, an incredibly general claim concerning fundamental concepts. Perhaps “metaphysical” claim would serve equally well here.

in everyday life, we are unable to imagine a world in sufficient disambiguated detail so as to confirm it as a p-world in this manner. Hence, where p is a far-reaching philosophical claim, we are unable to conceive of p. Thus, we are not justified in asserting the possibility of far-reaching philosophical claims.

Van Inwagen does not apply this argument directly to the usual gallery of famous philosophical possibility-claims. Rather, he seems to adopt the following strategy: taking a sample of two relatively conservative “far-out” possibility-claims - namely, that a naturally purple cow is possible and that transparent iron is possible - he argues that, in these two cases, Yablo-conceivability fails. He then appears to conclude that if Yablo-conceivability fails here, we can hardly expect to succeed in conceiving of propositions that concern subject matter even more mysterious than purple cows and transparent iron, such as God or the mind.

It would seem that we can imagine a naturally purple cow quite easily. Just picture one bounding around a meadow. However, this kind of imagining does not suffice for Yablo-conceivability, van Inwagen says [27, p.78]: such a scene is compatible with the falsity of the proposition that there is a naturally purple cow - the cow in our imagined field might just have been dyed. In order to imagine a situation that *verifies* the existence of a naturally purple cow, van Inwagen suggests, we would need to imagine a chemically possible purple pigment, such that this pigment is produced and properly placed in the cow's coat by structures that exist in some DNA that counts as cow DNA. Van Inwagen is doubtful that anyone has, or can, perform *this* imaginative exercise.

Likewise, van Inwagen thinks it isn't enough to, for instance, imagine a scientific conference in which leading scientists, to great acclaim, announce the existence of transparent iron to an applauding audience, while holding up a chunk of what looks like transparent glass [27, p.79]. This imagined situation is compatible with, for instance, the background scenario that transparent iron does not exist, but the scientific

community has somehow been deceived into thinking that it does. In order to verify the imagined existence of transparent iron, again an insight into the micro-structure of such a substance is required. And again, there is no evidence that anyone has, in fact, attempted this kind of detail when imagining this scenario.

3.3 Van Inwagen's argument formalized

Here, I think, is a reasonable, charitable first formulation of van Inwagen's argument (one that is certainly a great deal clearer than his own presentation of the argument):

van Inwagen's skeptical argument

- P1. One is justified in asserting the possibility of a proposition p only if someone has imagined a world that the imaginer takes to verify the proposition p .
- P2. Someone has imagined a world that the imaginer takes to verify proposition p only if someone has imagined a world in a sufficient amount of detail, relevant to p , so as to rule out the compatibility of the specified details of that world with $\neg p$.
- P3. If proposition p belongs to the class FP (where a proposition is a member of this class iff it is a philosophical proposition that is far-removed from everyday experience), then no-one has imagined a world in a sufficient amount of detail relevant to p .
- C. So, if proposition p belongs to class FP, then one is not justified in asserting the possibility of proposition p .

Let us take note of a few of the features of this argument. The first thing to note is that the argument is valid: if the premises are true, the conclusion will have to follow.

Let us now consider P1. What is particularly striking about P1 is that it is the *converse* of the weak c-p principle. The validity of the argument rests upon this. Indeed, the weak c-p principle occurs nowhere in the argument.

Let us consider P2. It seems safe to say that both Yablo and van Inwagen recognize this as a necessary condition of Yablo-conceivability. Do they, however, think that the converse of P2 also holds? That is, should P2 be reformulated as an *interpretation* and *clarification* of Yablo-conceivability? This is not so clear.

The issue of the correct interpretation/clarification of Yablo-conceivability is an important one (not only for van Inwagen's argument, but, of course, for Yablo's thesis too). I am certain that the reader will agree that, as Yablo presents it, his principle is somewhat vague. For all Yablo has said, it is understandable that one might not be sure what *exactly* is meant to be required of one, when performing imaginative acts, in assessing the possibility of a claim. Yablo makes it clear that he thinks the imagination is required and that only a specific kind of imagining will do, one involving "worlds" and the appearance of "verification". Yet what does this amount to? Understanding Yablo's conceivability properly is crucial to all three premises of van Inwagen's argument. This issue of interpretation means that the central question of this paper is somewhat complicated: in discussing whether Yablo's conceivability supports van Inwagen's skepticism, we cannot merely assume the truth of Yablo's conceivability and then try to discern whether van Inwagen's skepticism flows from it. It simply isn't all that clear what would be assumed in this case. Thus, in order to properly launch an analysis and defense of van Inwagen's argument, I will mainly devote chapter 4 of this paper to the development of what I see as a suitably clear and transparent theoretical framework by which to understand Yablo's principle. Only then can discussion around the truth of van Inwagen's premises begin in earnest.

What I *will* take as a background assumption, for my defense of van Inwagen, is the "spirit" of Yablo's conceivability: namely, the broad claim that the "conceiv-

ability” that is relevant to conceivability-possibility arguments is *imaginability*, but of a specific ilk. Whatever such imagining involves, it must account for both the success of the imagination and the sometime failure of the imagination in discerning possibilities, and it must make sense of the age-old gut-feeling that “clarity” and “distinctness” are crucial to such imagining.

Finally, let us consider P3. Like van Inwagen, I think that an inductive defense of P3 ought to suffice. The class FP will probably be difficult to define precisely and it certainly isn't feasible to attempt to run through every proposition that falls in this class. Demonstrating the truth of P3 for a small, yet representative, sample of FP-type proposition will be more than adequate. Unlike van Inwagen, however, I think it is important to pick genuinely significant philosophical propositions as our sample class. Thus, I will consider P3 as adequately defended if it can be shown to undermine the ghost, goblins, ghouls and sorcery arguments (which I introduced in an earlier section, section 2.4).

3.4 Objections to van Inwagen's argument

To provide impetus to continue our investigation, we now require some clear objections to van Inwagen's argument. If we are unable to crystallize any such worries, then we might as well accept victory on van Inwagen's behalf right now! Unfortunately for van Inwagen, reasonable objections to all three of his premises exist.

3.4.1 Objection to the first premise

Let us first consider an objection to P1. Assuming that Yablo's conceivability is sufficient for justified belief in possibility, why think that this form of conceivability is *necessary* for justification? This would certainly be beyond doubt if we were unable to come up with reasonable alternatives by which one can come to assert modal

knowledge. However, there *do* seem to be alternative methods which individuals sometimes cite as the justification for their modal beliefs. Those that readily spring to mind are the following: modal intuition, modal logic, the actuality principle and the similarity principle⁵.

Let me introduce each of the alternatives in turn. When someone cites modal intuition as the root of a belief in a possibility, what they mean, roughly, is that they can simply *see* that something is possible. Modal intuitions, like all intuitions, are real enough. However, they are only of interest as an alternative to other techniques of modal inquiry if they are not assumed to be a) based merely on habit, and therefore essentially foundationless, or b) based on extensive experience with using other modal techniques, and therefore parasitic on those techniques. Therefore, we will take modal intuition to refer to some kind of irreducible, primitive “inner vision” concerning what is possible and necessary.

Modal logic refers to any system of logic that introduces the \Box operator (it is necessary that...) and \Diamond operator (it is possible that...) to standard propositional or predicate logic, along with several axioms to regulate inferences for propositions containing these operators. A variety of such logics exist [16, 11]. There is some controversy as to which axioms defines a “correct” modal logic (correct not in the logical sense - that is, complete and consistent - but rather in the sense of which modal logic is a reliable guide concerning the possible and the actual).

The actuality principle is as follows: if p is actual, then p is possible. This principle enjoys little controversy.

Finally, let me introduce the similarity principle. Although this principle is seldom given a name, I think it is in wide use. Consider the following common answers to questions like “how do you know that the glass could break if dropped?” or “how do

⁵We can also acquire knowledge of possibility through the principle “if necessarily p, then possibly p”. Since I am avoiding discussing necessity, as far as possible, and how we come to be justified in asserting necessity, I will set this principle aside.

you know that you could leave this room by the door?" After a moment of puzzlement at being asked such seemingly inane questions, a layman is likely to respond "well, because things like that happen all the time", or perhaps, simply, "experience". That is, it seems that we often justify our modal beliefs (chiefly our most basic ones) according to our experience of the *actual* world. Yet this is not an application of the actuality principle. What is going on here? My best guess is that something like the following principle is being applied:

If two (or more) situations (or objects) are relevantly similar (to some degree), then the possibilities of those situations (or objects) are likely (to some degree) to be the same. The degree of likelihood of the similarity of possibilities is proportional to the degree of similarity of the situations (or objects).

This principle is utilized thus: in the past, I have left this room by the door. Thus, under those circumstances, it was possible for me to leave the room by the door (by the actuality principle). Now, the current situation in which I find myself is relevantly similar to the one which occurred in the past (the door is not locked, my body is functioning normally, and so on). Therefore, by the similarity principle, it is possible, under the current circumstances, that I leave the room by the door.

I won't defend this principle now, although I will return to it later in the paper. For now, I hope it is at least plausible that such a principle presents some kind of competition to Yablo conceivability.

In order to defend the truth of P1, then, one would have to show that for each one of these techniques, either that technique does *not* provide modal justification (appearances and opinions to the contrary), *or* that technique (perhaps implicitly) relies on Yablo-conceivability in order to provide the justification it does provide (and, thus, the technique assumes the truth of the c-p principle), *or* justification by that technique means it is guaranteed, or likely, that one is able to Yablo-conceive of the

proposition that has been justified (although the initial justification in this case is *independent* of Yablo-conceivability).

Let me immediately write off two of our alternatives as viable competitors to Yablo-conceivability. First, consider modal intuition. While an irreducible faculty or “vision” for determining what is possible may be attractive to some, I am not inclined to find intuitionism of this breed even vaguely plausible. Indeed, intuitionism of this brand faces the same problem as other forms of intuitionism: the impenetrable mystery that will descend upon our claims to knowledge if they rely purely on intuition. If intuitions disagree, how do we resolve the matter (and, clearly, there is disagreement amongst us as to what counts as possible)? By figuring out which of us has a more finely tuned sense of intuition? But how do we do that? Ultimately, if we wish modal inquiry to be transparent, we require techniques less subjective than reliance on mere intuition.

Second, consider modal logic. Modal logic remains embroiled in controversy as to which system is preferable. What is notable about this situation, however, is that commentators inevitably fall back on *conceivability*, in something like the Yablo sense, to defend their view that a certain rule of inference or modal axiom is ill-judged. For instance, Nathan Salmon argues that the characteristic S4 axiom schema $\Box p \supset \Box \Box p$ is untrue in many cases, by way of an imagined counter-example [18]. That is, Salmon can imagine a situation in which it is true that $\Box p$, yet false that $\Box \Box p$ and, thus, concludes that it is possible that $\Box p$ is true, while $\Box \Box p$ is false. What this seems to suggest is that modal logic is subject to the c-p principle, in the vein of Yablo: the c-p principle is used as an overriding principle for testing the adequacy of a modal logic, and is thus assumed to be true by the developers of modal logic.

Let us thus not concern ourselves further with modal intuition or modal logic. However, it is a trickier matter to discern if the actuality and similarity principles assume the truth of the c-p principle. It seems unlikely that the the actuality principle

assumes on the c-p principle. If the similarity principle assumes the truth of the c-p principle, it is not, on the face of it, clear how. What might sound more promising is the fact that possibility-claims that are justified by these principles are normally “basic” or everyday claims and hence, apparently, easily imaginable (it does not take much imagination to imagine a drinking glass breaking, or to imagine oneself leaving the room via an open door). However, does this kind of imaginability amount to Yablo-conceivability? As an objection to the second premise will show, it isn't clear that our “basic”, non-controversial modal knowledge is always Yablo-conceivable.

3.4.2 Objection to the second premise

One might accept the truth of the first of van Inwagen's premises and, yet still reject his argument by denying the truth of P2. This would amount to a rejection of the “reading” of Yablo's conceivability that gives rise to P2 (while still, perhaps, accepting the force of Yablo's basic idea).

Here is the objection, in the form of a *reductio* argument⁶:

Assume that both P1 and P2 of van Inwagen's argument hold. It follows that one is justified in asserting the possibility of a proposition p only if one imagines a world in a sufficient amount of detail, relevant to p , so as to rule out the compatibility of the specified details of that world with $\neg p$. However, there are cases of uncontroversially true modal statements where one has not, and probably will never, imagine a world of which that statement is a true description, in the kind of detail required. For instance, van Inwagen (and others) hold that we know that it is possible that John F. Kennedy should have died of a heart attack instead of by an assassin's bullet. However, no one has imagined such a scenario in

⁶This argument appears to be the thrust of Geirsson's central objection to van Inwagen's skepticism [8]

complete (relevant) detail, since constructing such a detailed world, with a detailed alternate history for an actual person, is, amongst other difficulties, simply too much work. A second example concerns even more mundane modal knowledge: I know that my dog could be lying in the sun right now and I seem to be able to picture this scene without difficulty. Nevertheless, I have not imagined this scene in the amount of detail required: nothing about what I have pictured is incompatible with the background scenario that my dog is *not* lying in the sun, but rather an exact copy of my dog is lying in the sun. Further, the background scenario stipulates that this clone was created by aliens immediately after the birth of my dog, upon which my dog was vaporized and replaced by his double. Now, what about my imagining my dog sleeping in the sun can rule out this background scenario? Nothing, it would seem, since I have no idea of the kinds of essential properties my dog holds (such as the identity of his parents) that would help to distinguish him from a double in this scenario. I could *stipulate* that I am imagining my dog, not a clone, but this kind of stipulation appears to miss the point of Yablo-conceivability. Thus, we have examples of possibility-claims that are clearly justified for a speaker to hold, but for which the consequent of P2 is false. Thus, by reductio, P2 is dropped, as the suspicious aspect of the assumption we began with.

The core of the objection is thus the complaint that P2 makes Yablo-conceivability too *demanding*. Notice that, if one were not attached to P1 but convinced that P2 is true, then the above argument serves equally well as an objection to P1. Thus, acceptance of the above reasoning can lead in two directions: either P1 is dropped, and it will have to be accepted that Yablo-conceivability is not the only guide to possibility, if it is a guide at all, or P2 is dropped, and a better reading, or update, of Yablo-conceivability will need to be offered (this is the route that Geirsson takes

in his paper, after presenting roughly the above argument [8])⁷.

3.4.3 Objection to the third premise

One might accept both P1 and P2 (that is, accept that Yablo-conceivability is necessary for justified belief in possibility, and accept that an adequate understanding of what Yablo-conceivability entails has been found) and still reject van Inwagen's argument on the strength of rejecting P3. One could do this by claiming that our sample of philosophical propositions are, in fact, assertions aside, Yablo-conceivable. Why might one think that ghosts, ghouls, goblins and sorcery are Yablo-conceivable? Because of the incredible lengths that some have gone to in imagining such things! One only has to look at a copy of the *Lord of the Rings*, or at the film version of the same, to witness the extent to which human beings have gone to create living, breathing worlds, rich in details and history, in which ghost, ghouls, goblins and sorcery exist (seemingly unequivocally). If this doesn't count as the conceivability relevant to possibility, the objector might say, then nothing does. If even more, unreasonable levels of detail are required, then we shall find ourselves back at the objection that Yablo-conceivability is too demanding, the objector concludes.

Thus, we have three incisive objections against van Inwagen's argument. These objections are obviously not unconnected. Indeed, the issue of formulating a clear and plausible reading of Yablo-conceivability bears on all three objections. Providing an analysis of Yablo's conceivability will be our next port of call.

⁷There is, in fact, a third option: to simply bite the bullet and deny that one is justified in asserting that JFK could have died of a heart attack or that one's dog could sleep in the sun. It seems unlikely though that anyone would be so attached to both P1 and P2 that they would go this route.

Chapter 4

The “Safe Explanation” Model

In this section, I aim to develop and motivate a reading of Yablo’s conceivability that will a) expand the insights provided by this version of conceivability into something closer to a comprehensive theory of modal epistemology, and b) give us a surer footing in evaluating objections to van Inwagen’s argument. In the process, I hope to take Yablo’s ideas further than what is offered in his original discussion.

I hope it is clear that Yablo’s account of conceivability *needs* some spelling out. I don’t think the account, as it stands, makes it entirely transparent what exactly is required of one’s imaginative powers in order to investigate what is possible. Furthermore, objections were raised against van Inwagen’s argument, in the last chapter, that bear upon the cogency of Yablo’s account and thus demand an answer from the friend of both Yablo-conceivability and van Inwagen’s skepticism.

To offer a glimpse as to what I will conclude at the end of this chapter, let me briefly summarize the account (of how we attain justified modal belief) that will emerge from Yablo’s insights and that I will develop throughout this chapter - what I call the “safe explanation” model of modal epistemology. Suppose one is wondering whether one is justified in asserting the possibility of proposition *p*. Now either “it is possible that *p*” is a basic modal proposition, or it isn’t¹. If “it is possible

¹One might have doubts that a sharp line can be drawn between basic and non-basic modal

that p ” is a basic modal proposition, then belief in the possibility of p is justified by either the actuality principle, or the similarity principle (in conjunction with the actuality principle). If, on the other hand, “it is possible that p ” is a non-basic modal proposition, then in order for one to be justified in asserting the possibility of p , one must imagine a consistent fictional world of which a set of propositions $p_1, p_2 \dots p_3$ is fictionally true, such that those propositions i) are such that one can deduce p (that is, make a deductive logical inference to the conclusion p) from those propositions and ii) are all less modally controversial than p . The combined degree of modal controversiality of $p_1, p_2 \dots p_3$, is inversely proportional to the degree to which one is justified in believing the possibility of p . I will refer to p_1, p_2, \dots, p_n as a “*modally safe*” explanation of p , one that gives a “safe” modal grounding to p .

4.1 An analysis of Yablo’s conceivability

Consider again Yablo’s preferred account of conceivability:

p is [philosophically] conceivable for me iff I can imagine a world that I take to verify p [29, p.29].

This account will, firstly, benefit from being phrased in a way that is as free of modal concepts as possible. Here is a revision:

p has been philosophically conceived by someone iff someone has imagined a world that the imaginer takes to verify p .

This account of conceivability is meant to be the best candidate for underpinning the weak c - p principle. Yet it remains somewhat mysterious what exactly is required of one, according to the above statement, when making possibility judgments. This

propositions, given the vagueness of van Inwagen’s, and my own (to this point), use of these terms. However, I will ultimately define these terms in a more precise fashion that I hope will make the distinction between “basic” and “non-basic” quite stark.

is no doubt due to the somewhat enigmatic reference to problematic notions such as “world” and “verification”. How exactly are we meant to understand these terms, in this context?

4.1.1 What is it to imagine a world?

Does Yablo instruct us to imagine a possible world or a fictional world, when conceiving in a manner relevant to discerning possibilities? Yablo appears to intend that one imagine a *possible* world: “According to [the account of conceivability], the task of conceiving p divides into two sub-tasks: imagining a possible world and satisfying oneself that p is true in it” [29, p.31]. This is an apparently worthwhile goal, since embedding proposition p in a possible world solves the problem of determining its possibility directly. What Yablo means by a possible world, it seems, is a world that is both complete and consistent (although asking for the completeness of a world is rather redundant - completeness is precisely what makes a world a world) [29, pp.28-29]. This is an admirable enough account of what a possible world is, in the sense that it is largely agreed that a comprehensive description of a possible world (the list of propositions true of that world) must be complete (for any proposition, it is either true of that world, or - exclusive or - the denial of that proposition is true of that world) and mutually consistent. Whether these conditions are not only necessary but also sufficient is another matter, but it must at least be agreed that if a world is complete and consistent, then this counts as positive evidence for that world being a possible world.

However, there is a problem if this is Yablo’s instruction. It is, in an important sense, beyond the capabilities of a human mind to imagine a possible world: it would require an omniscient being with an infinite mind to contemplate a “complete” world, in the sense relevant to possible worlds, and check that that world is consistent. The best we humans can do, in contrast, is to construct fictional worlds.

These correspond to a relatively small set of propositions (“fictional truths”). Thus, a comprehensive description of a fictional world can, at best, only approximate a comprehensive description of a possible world, and, hence, fictional worlds can only, at best, approximate possible worlds. So, perhaps, upon reflection of our limited capabilities, what Yablo *really* means to instruct us is not to imagine a possible world, but rather to imagine a *fictional world* that corresponds to or accurately approximates a possible world. In other words, his definition of conceivability ought to read:

p has been [philosophically] conceived by someone if[f] someone has imagined a consistent fictional world, which corresponds to a possible world (as a partial representation of that possible world), that is taken by the imaginer to verify p.

If this is indeed Yablo’s instruction, then again it runs into a crucial objection. An objector might complain that the above instruction takes for granted that we can tell when an imagined, fictional world corresponds to a possible world. If this is taken for granted, then the task of telling whether proposition p is possible is trivialized (since “I know that the fictional world w, that I have imagined, corresponds to possible world pw and p is a true description of w” implies that “possible world pw exists and p is a true description of pw” and “there exists a possible world of which p is a true description” can be paraphrased as “it is possible that p”), but this trivialization only masks the real problem: how to provide justification for the claim that we can tell when one of our fictional worlds in fact approximates a possible world! Thus, the objector concludes, Yablo has simply rephrased the central problem of linking conceivability/imaginability to possibility, without actually attempting to solve it².

This objection is misplaced, however. While we might not be able to conclude once and for all that any given fictional world corresponds to a possible world, there

²This appears to be the substance of the objection offered by Ernest Sosa in his article “Modal and Other A Priori Epistemology” [24].

are at least features that a fictional world can have that can give genuine weight to the hypothesis that a correspondence holds, in the absence of evidence to the contrary. First, if the fictional world has been constructed so that it is consistent (that is, the fictional truths of that fictional world are mutually consistent), this is some evidence in favour of it corresponding to a possible world, since it is generally accepted that a comprehensive description of a possible world will, at least, be consistent. Secondly, presuming that we have some uncontroversial modal knowledge, basic or otherwise, (let us not worry, for the moment, where this knowledge has itself come from), if a fictional world is also constructed so that all its fictional truths are uncontroversially possible, then this acts as further reason to posit a corresponding possible world. Let me define some terminology: I will say proposition p is *modally controversial* if we have no justification for asserting the possibility of p , or if we have justification for thinking that p is impossible. A proposition p is *modally uncontroversial* if it is not modally controversial. The second requirement thus suggests that we have reason for thinking fictional world fw corresponds to a possible world if it is constructed so that no modally controversial propositions are fictional truths. What is an example of a modally uncontroversial claim? Let us continue to suppose (I will give some reason for this later), that if “it is possible that p ” is a basic modal claim (for example, “it is possible that the furniture be re-arranged”), then p is modally uncontroversial. What are examples of modally controversial propositions? One obvious example is the denial of a generally accepted necessary truth, such as $2+2=4$. Another example is a proposition that makes a claim that we have little reason for believing holds in the actual world and that does not constitute a non-basic modal claim (for example, “it is possible that unicorns exist”). Thirdly, and in addition to the other two requirements, the more consistent, modally uncontroversial detail added to a fictional world, the greater the probability that it matches a possible world, since the closer the fictional world’s list of fictional truths will get to being “completed” (that is, the closer it will

get to having, for every proposition, either that proposition or its negation on the list)³.

At this point, our objector might again enter the fray, however. Surely it is with respect to the second requirement that Yablo goes wrong (the objector might say), since in imagining p as part of the fictional world, we are imagining a *controversial* modal proposition as part of the fictional landscape, which surely immediately disqualifies us from positing, with a clear conscience, a possible world that matches that fictional world.

I think this objection betrays a poor reading of Yablo’s conceivability, however (poor in the sense that a more useful reading is possible). To show this, and get a clearer idea of how I think we ought to understand Yablo’s conceivability, let us again consider the features of fictional worlds.

For any given fictional world and any given proposition, that proposition is either a fictional truth or a fictional falsehood of that world, or it is undecided if that proposition is true or false relative to that world. Of the fictional truths, some are *stipulated*. By a *stipulated fictional truth*, I mean a fictional truth which is true of the fictional world in virtue of a self-conscious decision of the author of that fictional world. For example, an author might self-consciously set the colour of the hat their lead character is wearing to be the colour red. “The colour of the hat is red” is then a stipulated fictional truth. Stipulated fictional truths are thus true of a fictional world by virtue of a knowing decision by the author and therefore *it cannot be the case that that proposition is not fictionally true of that fictional world*. It also cannot be the case that the author can be mistaken as to the truth value of that particular proposition. It might be the case that stipulating that certain propositions hold for a fictional world render that world inconsistent, but so be it: some fictional worlds simply are inconsistent.

³The reason this third requirement must be in addition to the others, is that more detail being added is not necessarily a good thing if it is not consistent and modally uncontroversial.

One can separate the fictional truths of a fictional world into *foreground truths* and *background truths*. A proposition p is a foreground fictional truth just in case the circumstances described by p have been represented by the author in detail. “HAL9000 asked Dave what he is doing” is a foreground fictional truth in the fictional world of *2001: A Space Odyssey*. A fictional truth p is a background truth if that proposition is fictionally true of the world, but p has not been represented by the author in detail. For instance, in the *Spiderman* comic books, Peter Parker has unusual DNA that grants him spider-like abilities. However, while the proposition “Peter Parker’s DNA has certain unique properties” is undoubtedly true of the *Spiderman* stories (denial of this proposition amounts to a confusion as to the facts of this fictional world), no author of a *Spiderman* yarn has imagined this DNA in detail and so this is a background truth of this fictional world. Both foreground truths and background truths can be stipulated, since both can be set through a self-conscious decision of the author.

Now, the objection we were discussing appears to turn on the ease with which we can stipulate either impossible propositions, or propositions whose possibility can legitimately be questioned, in fictional worlds (even when those fictional worlds appear consistent). This is easy to see in the case of background truths. Indeed, it seems almost unlimited what background truths an author can stipulate for their fictional world. I can quite easily concoct a fictional story that revolves around a well-known mathematical theorem being false in that fictional world (the story is a thriller: a clandestine society of mathematicians, with assistance from the CIA, has for years been suppressing the falsity of the theorem for their own nefarious purposes). If, however, the theorem is, in fact, true, then a crucial background truth in the fictional world is impossible. However, one might hope that impossibilities cannot be stipulated to be foreground truths, since there appear to be many impossibilities that we cannot imagine as part of detailed situations (I cannot imagine a situation in detail

where two melons and two melons add up to five melons). Does this present a means for escape from the objection? Unfortunately not. Despite some seeming restrictions, many impossibilities can be imagined as part of reasonably detailed situations. I can distinctly imagine a situation in which I shake hands with myself (I can paint a very detailed picture of the imagined scene, if required). I can also distinctly imagine a situation in which George W. Bush’s skin is peeled back and it is revealed that he is not a human being, but rather a Soviet automaton. I can go on. The objection, then, amounts to this: Yablo asks us to stipulate of a fictional world that a modally controversial proposition p holds for that world, and then also assumes that we have some way of knowing that that fictional world corresponds to a possible world. However, this last assumption is an unjustified one, since stipulating that a modally controversial proposition holds for a fictional world (whether it be as a background or a foreground truth) deprives us of whatever support we have for a correspondence between our fictional world and a possible world, since there are clear cases in which impossibilities are stipulated as holding for fictional worlds.

In reply to the objector, it seems to me that this would only be a problem if Yablo-conceivability *were* asking us, or at least allowing room for us, to directly specify, or stipulate, that p holds of the imagined fictional world. However, Yablo-conceivability does not suggest that we simply stipulate p , as a foreground truth or otherwise, to be fictionally true of our imagined world - the instruction is that we attempt to imagine the world so that p is *verified* as being true of that world. The idea appears to be that the world one has constructed ought to *lead* one to assert the truth of p in that world (without having committed to the fictional truth of p beforehand), by way of the other fictional truths of that world (which themselves may have been stipulated). Thus, p will be a fictional truth that is not stipulated. Can there be fictional truths that aren’t stipulated, given that it is accepted that the fictional world is constructed by its author? This appears to be the case. For one

thing, when the author imagines a scenario within a fictional world, the author can be surprised by how exactly their imagination “fills out” the details of the situation. I’m not certain that such “surprising” details can qualify as stipulated. More importantly, if the author’s goal is to have a consistent fictional world and, by some process of reasoning, it is shown that the stipulated truths of a fictional world indicate further truths about that world, then these surely count as fictional truths which are not stipulated (the author of a novel might be surprised to have a lapse in continuity pointed out to them by a reader, for instance). With these options in mind, how should we understand “verification” within a fictional world?

4.1.2 What is it to verify a proposition within an imagined world?

The insight that Yablo has latched on, it seems to me, is thus to insist that we must not stipulate the truth of p in our fictional world. Rather, the other, less modally controversial facts of our fictional world ought to lead one to accept p as being true of that world. The fictional world must *verify* p in some way.

“Verification” is a word that is used, on different occasions, to refer to two quite different kinds of epistemic support. On the one hand, when it is said that “it has been verified that p ”, it is sometimes meant that *evidence* has been found in support of proposition p . Let us say, to be more precise, that statement e is *evidence* for p just in case e is logically inferable from p (along with other factors, perhaps), and e is in fact true. Hence, e *verifies* p , in the current sense, just in case e is evidence for p . Verification of p , in this sense, thus amounts, at best, to accruing *fallible* support for e . When spoken of in this manner, verification is contrasted with confirmation, where p has been confirmed just in case support for p has been accrued that *guarantees* the truth of p . Is this the kind of verification by which Yablo’s conceivability is most fruitfully understood? If this is the case, it would appear that the directive to imagine

a world that verifies p , translates to a directive to imagine a world in which there is merely evidence, but no proof, for p . We must imagine, in other words, acquiring evidence for p . In light of what was said above, this must obviously be done *without* stipulating the truth of p : we want to be *convinced* of the fictional truth of p .

This way of taking “verification” will not do in the context of imagining fictional worlds. Imagined evidence is simply too cheap: I can imagine plenty of evidence that Clarabell is a naturally purple cow: Clarabell’s purple tone never fades; Clarabell was born of two equally purple cows; and so on. None of this seems the slightest bit persuasive however, with regards to convincing one of the possibility of a naturally purple cow. This lack of persuasiveness is no doubt due to the fact that imagined evidence is *made up*, and it seems crucial to the persuasiveness of evidence that it be *discovered*.

The second way of taking “verification” is more hopeful. Sometimes the word “verification” is used in the same manner as “confirmation” was used in contrast to our last reading of “verification”: to verify p is to find *conclusive* reasons for believing that p is the case. In other words, p_1, p_2, \dots, p_n verify p when p_1, p_2, \dots, p_n guarantees the truth of p - that is, p_1, p_2, \dots, p_n jointly entail p . What I mean by entailment is that it is logically impossible for p_1, p_2, \dots, p_n to all be true while p is false. In this case, the directive to imagine a world that verifies p , amounts to a directive to imagine a world such that p_1, p_2, \dots, p_n are all true, where p_1, p_2, \dots, p_n entail p . This immediately sounds more promising: our goal is to have a world of which p is true while avoiding directly stipulating the truth of p and this can be accomplished by directly stipulating the truth of propositions that *entail* p .

At this point, the discussion has yielded the following clarification of Yablo’s conceivability:

p has been philosophically conceived by someone iff someone has imagined a consistent fictional world of which some set of propositions $p_1, p_2, \dots,$

p_1, p_2, \dots, p_n are true, such that p_1, p_2, \dots, p_n jointly entail p .

I remind the reader that “consistent” here simply means that the fictional truths of the fictional world are mutually consistent. Now, since the consistency of the fictional world and the truth of p_1, p_2, \dots, p_n together ensure that that world is incompatible with the truth of the denial of p , one may be tempted to suggest that the most fruitful reading of Yablo’s conceivability overlaps with the second premise of van Inwagen’s skeptical argument, as I presented the argument earlier:

p has been philosophically conceived by someone iff someone has imagined a world in a sufficient amount of detail, relevant to p , so as to rule out the compatibility of the specified details of that world with $\neg p$.

However, we need to be cautious here. Two issues remain with this “entailment” reading of Yablo-conceivability. The first is this: we would like it to be the case, for the philosophical conceivability of p , that the details of the fictional world will *lead* one to see that p is fictionally true of that world, given that p has not been stipulated (and so it is not a given that p is fictionally true of that world). What we would like, it seems, is that p be *inferred* from the details of the world (some set of fictional truths p_1, p_2, \dots, p_n). However, it seems that the mere entailment of p from fictional truths p_1, p_2, \dots, p_n is not enough to ensure that p can be inferred from p_1, p_2, \dots, p_n . To illustrate this, let p_1 be “the USA invaded Iraq in 2003” and let p be “ $2+2=4$ ”. Now, since “ $2+2=4$ ” can never be false (as a necessary truth), it is not logically possible that p_1 be true while p be false, and, hence, p_1 entails p . However, I would hardly say that “ $2+2=4$ ” can be *inferred* from “the USA invaded Iraq in 2003”, in the sense that I do not think it is possible that if someone were uncertain of the truth status of “ $2+2=4$ ” that they could justifiably be convinced of it purely on the basis of knowing the truth status of “the USA invaded Iraq in 2003”! What this amounts to is that the above account of philosophical conceivability is flawed: as it stands, it claims that one

has philosophically conceived of “ $2+2=4$ ” when one has imagined a scenario, with no inconsistencies, in which the USA invaded Iraq in 2003. This does not seem to be an instance of conceiving of “ $2+2=4$ ”, in any sense of “conceiving”. This situation can be rectified by amending the account so that “ p_1, p_2, \dots, p_n jointly entail p ” is replaced by “ p is deduced from p_1, p_2, \dots, p_n by the imaginer”, where p is *deduced* from p_1, p_2, \dots, p_n by person x just in case i) p_1, p_2, \dots, p_n jointly entail p and ii) p is justifiably inferred⁴ from p_1, p_2, \dots, p_n by person x . Thus, we now have:

p has been philosophically conceived by someone iff someone has imagined a consistent fictional world of which some set of propositions p_1, p_2, \dots, p_n are fictionally true and p is deduced from p_1, p_2, \dots, p_n by the imaginer.

Another important detail has been neglected that, when taken into account, shows the inadequacy of the above reading of Yablo-conceivability. It was agreed earlier that if a fictional world is a) consistent, b) reasonably detailed and c) has no fictional truths which are modally controversial, then we have good reason (or at least the best reasons we can hope to find) to suppose, in the absence of evidence to the contrary, that that fictional world corresponds to a possible world, as a partial representation of that possible world. Now, if any of p_1, p_2, \dots, p_n are modally controversial, or at least as controversial or more controversial than p itself, then there is little reason to take one’s fictional world as indicating the possibility of p . Here’s an example to illustrate this requirement: it is obviously futile to attempt to show the possibility of somebody running the three-minute mile by imagining a fictional world of which it is true that God exists, it is true that God decrees that a human being with the relevant running ability pops into existence at time n and it is true that God compels this runner to run the three-minute mile at time $n+1$ (given God’s omnipotence, these propositions entail the truth of the proposition that the three-minute mile has been

⁴A justified inference occurs when one accepts the truth of a certain proposition on the strength of other propositions being true, where this acceptance is the result of reasoning according to truth-preserving inference rules. A formal derivation, for instance, counts as a justified inference.

run in this fictional world). The futility of this exercise, quite clearly, lies in the fact that the possibility of God existing is itself a controversial modal proposition, and, I would say, more so than the possibility of somebody running the three-minute mile.

Here then is the revisionary and clarificatory account of philosophical conceivability that emerges when taking a hard look at Yablo’s conceivability:

p is philosophically conceived by someone iff someone imagines a consistent, reasonably detailed fictional world of which some set of propositions p_1, p_2, \dots, p_n are fictionally true, such that i) p is deduced from p_1, p_2, \dots, p_n by the imaginer and ii) p_1, p_2, \dots, p_n are all less modally controversial than p.

Notice that this account means that if p has been philosophically conceived by me, then I have imagined a world in a sufficient amount of detail, relevant to p, so as to rule out the compatibility of the specified details of that world with $\neg p$. Hence, this account supports premise 2 of van Inwagen’s argument. It explicitly denies, however, that this necessary condition is also sufficient.

4.2 The “safe explanation” model of modal epistemology

Let me introduce some terminology. If both i) p has been deduced from p_1, p_2, \dots, p_n and ii) p_1, p_2, \dots, p_n are all less modally controversial than p, then I will refer to p_1, p_2, \dots, p_n as jointly constituting a “*modally safe*” explanation of p. I think the use of the word “explanation” is justified in this context. A rough account as to what an explanation is that it is some set of statements such that i) the thing to be explained can be logically inferred from those statements and ii) the acceptance of those statements removes the problematic or surprising character of the thing to

be explained. These conditions are, respectively, fulfilled by i) p has been deduced from p_1, p_2, \dots, p_n and ii) p_1, p_2, \dots, p_n are all less modally controversial than p . While I think the use of “explanation” is under the circumstances natural, I will not attempt to find a place for “modally safe” explanations within any detailed theory of “scientific” explanation, or anything along those lines.

I will thus refer to the account of modal epistemology that emerges from my reading of Yablo-conceivability as the “safe explanation” model of modal epistemology. According to this account: in order to be justified in asserting that proposition p is possible (where p 's possibility is controversial, relative to some set of uncontroversial possibility-claims), one must imagine a consistent fictional world, in as much detail as is reasonable (the more detail, the better), so that not only is p true of that world, but a “modally safe” explanation as to *how* p is true of that world (propositions p_1, p_2, \dots, p_n) can be gleaned from the details of the fictional world. The point of these restrictions on one's imagining is to make sure that the best reasons we can have to assert this fictional world as corresponding to a possible world, are upheld by the fictional world: the fictional truths of the fictional world are mutually consistent, a fair amount of detail has been added to the world and no modally controversial statements are true of that world - or, at least, no statements more doubtful than p are true of that world. In particular, one is not allowed to *merely* stipulate the truth of p . The more consistent detail added to the world, the more complete the explanation for p is, and the less controversial the propositions that constitute this explanation, the more justification one has for asserting the possibility of p .

Notice, first, that this account leaves room for there to be more to what makes a world possible than just consistency and completeness. For instance, we appear to know the truth of some statements of necessity through conceptual or semantic considerations, or other means. It is necessary that the angles of a triangle add up to 180 degrees; it is necessary that all bachelors are unmarried men; it is necessary that

a particular human being have originated from certain genetic material, and so on. Any claim that denies such necessity statements is clearly modally controversial. For instance, it is dubious, in the interests of concluding possibility, to explain that, in her fantasy world, a little girl is a princess because she is the daughter of the (fictional) king and queen in that world (we have good grounds for denying that this could be the case, and, hence, it is modally controversial).

Second, notice that the above account of modal epistemology leaves room for varying degrees of justification, and offers some insight as to why modal error can occur. A fictional world might fulfill all the given criteria and yet still fail to match a possible world, for reasons that are not immediately transparent to author of that world. Sometimes filling out further details of the world, or reflecting on the implications of already stipulated details, can reveal an inconsistency, for instance.

Before we can use this account to tackle van Inwagen’s skepticism, however, it needs to be fleshed out a little more. In the next few sections, I will make the account more comprehensive, dealing with a few worries along the way.

4.2.1 Believability and the blueprint paradigm

I have offered some principled philosophical reasons for accepting the above reading of Yablo-conceivability as the most suitable in the context of conceivability-possibility arguments. However, doubts may arise if there is little evidence that human beings do, in fact, use this kind of conceivability when making judgments about possibility. What is needed then are some examples that illustrate the *intuitive* plausibility of the “safe explanation” model - examples that establish a match between the above theory and how we reason about possibility on the ground.

I think there is ample evidence that we use a “safe explanation” model to make possibility-conclusions. Consider what it takes for one to be convinced of the “plausibility” or “believability” of fictional story-telling, where some of the circumstances

described in the story are considered modally controversial. Generally, a fictional story (its characters, plot and so on) is praised as believable when it is agreed that the story is filled with *realistic* details (and by realistic, one normally means details that match, to some degree, what we think we know about the actual world) that add weight to the more far-out details of the story. The more the author succeeds in filling out such details, the more one gets the sense that this story is no mere fanciful yarn, but an account of how things *could* have been. A striking example comes, once again, from the worlds of comic books. No major superhero is introduced without an “origin-story”. The origin-story *explains* why that character has the powers that he or she does. Such explanation is seen as vital to the integrity of the story, since it would be viewed by readers as arbitrary and fanciful to have people fly around, shoot rays from their hands or move objects with their minds, without *any* kind of explanation. These origin stories, as far as they are able, attempt to root these otherworldly abilities in “reality” (in particular, they attempt to make them compatible with the laws of logic and the actual laws of nature): powers are explained by way of advanced new technology, genetic mutations, gamma radiation, or, if the writers are feeling desperate, alien visitation. I won’t say that such explanations are always successful in making characters more believable (there is certainly no attempt by the authors of comic books to meet rigorous philosophical standards of justification), but the intention in offering such explanation is quite clear.

I am claiming that when we ask for an explanation as to the possibility of certain circumstances in a fictional world, what we are really after, upon reflection, and what an obliging author will attempt to provide in order to appease his audience, is a “modally safe” explanation. If, however, there is a coincidence between the demand for “believability”/“realism” in fiction and the theory I proposed earlier, what might p_1, p_2, \dots, p_n and p (from my revised statement of philosophical conceivability) be in, for instance, a superhero case? Here is an example:

- p1: There exists a human DNA structure, X, that is sufficiently close to that of a spider, so that a human being with that DNA structure has spider-like abilities (can walk on walls, and so on).
- p2: If a human being is bitten by a certain kind of radio-active spider, then, by some causal process, that human being’s DNA will be altered so that it has structure X.
- p3: Peter Parker was bitten by a radio-active spider.
- p: (Therefore:) Peter Parker has spider-like abilities.

This, I think, sums up the origin-story of Peter Parker/Spiderman. Phrased this precisely, it is no surprise that this origin-story fails to persuade one that a spider-man could exist: p1 through p3 are all as (or almost as) modally controversial as p. That this is a poor attempt at a “modally safe” explanation is beside the point, however: the implicit appeal to the above deductive reasoning, and the hope that an appeal to the scientific notion of “DNA” will add some modal weight to p1 through p3, is, I think, clearly an attempt at using the “safe explanation” framework.

More pedestrian examples exist in regular fiction: if it is difficult to believe that a human being could be quite as cold, calculating, ruthless and murderous as the villain in a thriller, then the author can, for example, provide a background story showing how family trauma early in the criminal’s life, of a type familiar to everyday life, led him to this state of mind. The “modally safe” explanation offered in such a case might, roughly, be something like this:

- p1: If a person has a naturally sensitive disposition, lacks moral guidance in their formative years and is exposed to (some specified kind of) severe trauma as a child, then that person, in adulthood, will drastically fail to recognize other human beings as being at all valuable.

p2: Hannibal Lecter experienced childhood trauma, a lack of moral guidance and has a naturally sensitive disposition.

p: (Therefore) Hannibal Lecter, as an adult, cares disturbingly little for other human beings.

In practice, of course, authors attempt far more nuanced explanations than the above, in the interest of believability. Notice that this argument fares better as a “modally safe” explanation than the spider-man example: p1 above has some support from actual world experience (which translates into support for its possibility). As for p2, this proposition itself seems open to “modally safe” support, if the author constructs the events of Lecter’s life accordingly.

A second kind of example, of where we implicitly use a “safe explanation” model in daily life, comes from the manner in which we generally decide that it is possible that man-made objects with certain properties could exist. Quite a number of the machines that men have built have, for non-experts, a miraculous quality about them: it is hardly a trivial matter that powerful microprocessors, cellphones and plasma screen TV’s exist. If such items were not as ubiquitous as they in fact nowadays are, it would be understandable that one would be wary to commit to their possibility. What would it take to convince one of the possibility of this kind of technology? A direct encounter would suffice, no doubt. One need *not* see a machine in the flesh to be convinced that it is possible, though: one, it must be agreed, would be equally satisfied to see, and understand, detailed *blueprints* for that machine. What is the function of these blueprints? They serve not only as a *plan* for how to realize a machine, but they also act as an *explanation* (in particular, when the information contained in the blue-print is translated into a series of propositions), in terms of uncontroversial elements (such as the assemblage of certain parts, or mechanical principles derived from experience), as to how such a machine could be. What is more, the more detailed the blueprints, the more weight one is likely to give the hypothesis that such a machine

is possible. This, it seems to me, is a paradigm case of the “safe explanation” model at work. Before the Wright Brothers, the possibility of a man-made object with the ability to fly long distances, made of wood and steel and carrying passengers, was a legitimately controversial modal proposal. The Wright brothers, however, would no doubt have had some success winning over a naysayer by providing the broad aerodynamic principles behind such a craft. They would have had even further success with a clear, detailed blueprint for the construction of an aeroplane. Similarly, one may legitimately wonder if a teleporter could exist, such as the one depicted on *Star Trek*. What would convince one that such a thing is possible? It certainly would go a long way to convincing the uncertain if an explanation could be provided, in terms of broad principles, as to how such a machine could operate. Even more compelling would be a blueprint of the machine. What would be most compelling of all (other than actually encountering such a machine) would be to see a blueprint that makes it unequivocally clear how one could build such a thing. Again, it seems to me that what fuels this growing conviction is the “safe explanation” model in action.

To again be a little more precise, consider the following example: suppose a renowned architect designs, in rough, a building called Extravagant Towers. In hopes of achieving lasting fame, this architect has designed Extravagant Towers to be, if built, the tallest building in the world at x feet and also to have the shape, in profile, of a ship’s sail. Suppose that there follows some dispute amongst engineers whether it is possible, given the materials at hand and the laws of nature, to realize the architect’s vision. However, after years in the doldrums, a dapper young engineer revives the project by presenting a detailed blueprint for the construction of the building. Even before the construction begins, it is realized by all that Extravagant Towers is possible, by way of imagining a fictional world in which Extravagant Towers stands, and p_1 and p_2 , as in the following “modally safe” explanation, hold:

p_1 : If a building is constructed using a steel frame of shape x , load-bearing materials

y located at strategic positions z, . . . and so on, as detailed in the accompanying blueprints, then such a building is not in danger of collapsing, given the actual laws of nature.

p2: Extravagant Towers is constructed using the principles outlined in the antecedent of the proposition above.

p: (Therefore) Extravagant Towers stands.

4.2.2 Explanations, infinite regress and circularity

There are two nagging worries about the “safe explanation” model, as it stands: the account seems to be in danger of either circularity, or an infinite regress. On the one hand, circularity might be a problem, since the method for discerning the possibility of a proposition, according to the account, requires one to *already* hold some justified modal beliefs, as one is required to construct a world using modally uncontroversial statements and avoiding modally controversial statements (since “controversial” is here synonymous with “unjustified” and “uncontroversial” is synonymous with “justified”). On the other hand, infinite regress might be an issue, since if *all* statements of possibility are to be justified using the above technique, then one might find oneself having to provide endless explanations: proposition p1 can only be justified/explained by (amongst others) proposition p2; but proposition p2 itself needs to be justified/explained, which requires proposition p3; now, proposition p3 needs to be explained/justified too, by proposition p4; and so on.

These two worries are not unconnected. A solution to the infinite regress problem will also solve the circularity problem. In order for an infinite regress to be halted, the explanations need to halt somewhere. Thus, there need to be propositions whose possibility does *not* need to be justified by imagining a consistent fictional world whose details provide a “modally safe” explanation for that proposition. Hence, there need

to be “foundational” uncontroversial possibility claims for the “safe explanation” model to work. Furthermore, if such foundational possibility claims are explicitly acknowledged, then the fear of a vicious circularity is abated. A system of explanations will suffice to justify non-foundational possibility claims, the explanations serving to link the non-foundational claims to the bed-rock of foundational claims.

All that remains is to give some account as to what these foundational claims are, and what justifies them. First, notice that the idea of two distinct categories of possibility-claims (in this case, foundational and non-foundational) is hardly an unfamiliar one. I think we tend to see a difference between our everyday modal claims, the kinds of claims that allow for normal living and that we generally assert and believe without question, and the more ambitious modal claims brought to our attention by the arts (particularly in genres such as science fiction or fantasy), philosophy or science. We may call the former “basic” modal claims (examples being: “it is possible that my living room furniture be rearranged”, “it is possible that I got out of bed late this morning”, “it is possible that this window pane should shatter”), while the latter are “non-basic” (“it is possible that God exists”, “it is possible that alien life exists on another planet”, “it could have been the case that the ice age didn’t occur and dinosaurs still roam the earth”). What exactly is the distinction being drawn here? In the first place, basic modal claims are somewhat sacrosanct (somewhat unsurprising considering how important they are to normal life): for instance, a theory of modal epistemology or modal metaphysics is likely to be viewed with suspicion if it suggests that we are not justified in believing basic modal claims. Non-basic modal claims, on the other hand, are, apparently, open to revision, and confident affirmation of the truth of one is liable to be met with suspicion. In the second place, an obvious distinction between the two types of claim is their “distance” from actual experience. Basic possibility-claims are undoubtedly closer to our everyday experience of the actual world. A world in which I got out of bed late this morning is hardly very different

to the actual world. What is more, things like that *in fact* happen quite frequently. Non-basic claims, on the other hand, are typically far-removed from the actual state of things.

We thus have an intuitive idea of a distinction between possibility-claims that are basic and non-basic, foundational and non-foundational, uncontroversial and controversial. The suggestion is thus that we are entitled to believe the basic possibility-claims, and can then construct justificatory explanations for non-basic claims, true for a fictional world, that link those states of affairs back to the basic claims which hold for that fictional world. The justification we have for believing the basic claims is thereby “transferred” to the non-basic claims.

Can a more in-depth account be given for our access to basic modal claims, other than simply stating that they are sacrosanct? Here are two options.

Option 1: Perhaps there are events in the history of the world that count as “branching points” - at these points, it is simply and irreducibly the case that the world’s history had the opportunity to branch off in several directions. A basic modal claim describes the immediate outcome of the world having traveled down one branch rather than another.

What could these “branching” events be? One option is that nature is non-deterministic. If the laws of nature are probabilistic, for instance, then branching events must occur: if a law of nature has the form “if event e occurs, then either p_1 (with probability x) or p_2 (with probability y)”, then event e is a branching event, and p_1 could have followed e , or p_2 could have followed e . The problem, however, with locating branching points in the realm of physics is that this hardly explains our everyday basic possibility-claims, which are known by physicist and layman alike.

Perhaps a more promising avenue is to explore the idea that *free choice* creates branching points in the history of the world. That is, whenever a genuinely free agent genuinely makes a choice, then the history of the world goes off along one of several

possible paths. Thus, it is true that I could go for a cup of tea right now, and it is true that I can remain here right now and it is true that I could go for a walk right now, because I am in a position to make a free choice right at this very moment. The existence of autonomous agents *creates* possibility for the world, by virtue of their freedom. Indeed, this kind of modal picture might be exactly what we need to understand the nature of freedom.

The problem with this picture, however, is that it ties an account of basic modal knowledge very closely to philosophical puzzles concerning freedom. For instance, some philosophers (hard determinists) hold that we aren’t free at all. In conjunction with the view that free choice grounds possibility, this outlook results in the view that possibility is an illusion. Other philosophers (compatibilists) hold that we are free, but that the world is nevertheless deterministic, since making a free choice is simply a matter of acting on one’s reasons, although one’s reasons are themselves determinate at any given moment. This view is, apparently, incompatible with the idea that free choice grounds possibility, since when a free choice is made on the compatibilist’s view, the world does not branch, it follows a determinate path. Only the libertarian can fruitfully pursue the above suggestion.

I would prefer to avoid the debate concerning the nature of freedom. Thus, I will drop this line of investigation and consider a second option as to the nature of our basic modal claims.

Option 2: For our second option, I suggest the following: given both our penchant to cite *experience* as the basis for our basic modal claims, and the undeniable “closeness” of basic modal claims to the actual world, it may prove fruitful to consider if experience of the actual world does indeed back up such claims. In section 3.4.1, I mentioned two principles that may set the stage for the suggestion that experience bolsters modal claims: the actuality principle and the similarity principle.

The actuality principle is the following: if p is actual, then p is possible. This

principle is basically beyond reproach. It is also a fruitful source of modal knowledge: what better way to convince someone that a certain state of affairs, or kind of object, is possible, than to actually *produce* them for observation? Indeed, it is hard to imagine better justification for a possibility-claim. A reliable way to convince someone that traveling to the moon in a spacecraft is possible, is to *actually* travel to the moon in a spacecraft.

Nevertheless, the actuality principle, in isolation, is obviously not satisfactory for accounting for all our basic modal knowledge. I know I could go for a cup of tea right now; but, I haven't; so the actuality principle cannot ground this exceedingly trivial possibility-claim. What is more, it will not do for grounding non-basic modal claims, which depart widely from actuality.

Thus, I suggest we also bring the similarity principle into the picture. Here is the principle, as I stated it in section 3.4.1:

If two (or more) situations (or objects) are relevantly similar (to some degree), then the possibilities of those situations (or objects) are likely (to some degree) to be the same. The degree of likelihood of the similarity of possibilities is proportional to the degree of similarity of the situations (or objects).

Again, here is an example of how the principle is utilized: in the past, I have dropped and smashed drinking glasses on the floor. I have a distinct memory of such events (and thus find it easy to imagine them happening again - hence, this is a non-actual scenario I recognize and take seriously). By the actuality principle, I know that it is possible that those past drinking glasses should drop and shatter. Now, if I am given a drinking glass, I immediately judge it to be relevantly similar to the drinking glasses I have dealt with in the past, in particular the ones which I dropped and broke. Thus, by the similarity principle, I conclude that it is possible that *this* drinking glass should break if dropped - and I make an effort to be careful with it.

The similarity principle, importantly, justifies (basic) possibilities that are non-actual. For instance, it is certainly not modally controversial to claim, as part of a fictional story, that a character in that story enjoys to smoke. Plenty of people actually enjoy smoking, and the author of the fiction may not have given us any reason to doubt the similarity between the fictional character and those actual smokers.

4.2.3 Defending the similarity principle

The similarity principle is admittedly not foolproof. While it often provides powerful justification for asserting a possibility, it can sometimes lead us astray. Chiefly, this can occur because we are not aware of a relevant difference between two situations or objects that are being compared (one might conclude that it is possible that a certain man could fall victim to a heart-attack, since one knows of no relevant differences between him and the many people who have fallen victim to a heart-attack. However, it turns out that this fellow is a fitness and health fanatic, with a history of strong hearts in his family). However, I think there is a reasonable inductive argument to justify confidence in the similarity principle.

Thanks to the actuality principle, there is a conclusive test which confirms the possibility of a proposition. To confirm that p is possible, make it so that p is actual. Obviously, there are plenty of possibility-claims which, in principle, cannot be tested in this way. Nevertheless, this test is a useful gauge for the similarity principle’s effectiveness: there are many predictions of possibility which the similarity principle makes which can be confirmed by the test. For instance, the similarity principle concludes that it is possible that I can leave the room right now. I can test this by actually leaving the room. Examples like this are, clearly, innumerable and the similarity principle, I am sure it is agreed, tends to fair very well in the face of such tests. What this amounts to is that significant evidence exists for the truth of the similarity principle. Thus, one may conclude inductively that the similarity principle

is true.

Thus, the similarity principle can be justified by way of the principle of induction conjoined with the actuality principle. Indeed, I think the connection between the similarity principle and the principle of induction (or the principle of the uniformity of nature), is a close one. Inductive conclusions seem to carry modal information: if it is concluded, inductively, that all metals expand when heated, then the next time I pick up a piece of metal, I am justified in asserting the non-actual possibility claim “it is possible that this piece of metal will expand when heated”. I won’t explore the connection between the similarity principle and principle of induction any further here, but, suffice to say, I think their respective statuses are closely intertwined. Considering the ubiquity of inductive arguments, despite philosophical controversy around them, this counts in the favour of the similarity principle.

4.3 Conclusion

We have now developed a fuller account of modal epistemology. Let us give a more precise definition for a “basic” modal claim: a *basic* possibility-claim is one that is (directly) justified by either the actuality principle, or the similarity principle (in conjunction with the actuality principle). A *non-basic* possibility-claim is a possibility-claim that cannot be directly justified by the actuality and similarity principles (that is, by experience). Thus, we have the following as a final formulation of the “safe explanation” model:

A basic possibility-claim is justified iff it is supported by either the actuality principle, or the similarity principle in conjunction with the actuality principle.

A non-basic possibility-claim “it is possible that p” is justified iff someone has imagined a consistent fictional world such that propositions p_1, p_2, \dots ,

p_n hold for that world, where i) p is deduced from p_1, p_2, \dots, p_n by the imaginer and ii) either $\diamond p_k$ ($1 \leq k \leq n$) is a basic possibility-claim, or $\diamond p_k$ is a non-basic possibility-claim that has been justified to some acceptable degree.

Thus, we end this section with a recursive account as to how one can be justified in asserting possibility-claims. The second statement, that for non-basic possibility-claims, obviously encapsulates a version of the weak c-p principle. The kind of conceivability touted here I will continue to refer to as “Yablo-conceivability”, in order to emphasize its roots. Now let us consider this account in relation to van Inwagen’s skeptical argument.

Chapter 5

Skepticism Reconsidered

With a more sophisticated take on Yablo-conceivability at hand, let us return to van Inwagen's skeptical argument:

van Inwagen's skeptical argument

- P1. One is justified in asserting the possibility of a proposition p only if someone has imagined a world that the imaginer takes to verify the proposition p .
- P2. Someone has imagined a world that the imaginer takes to verify proposition p only if someone has imagined a world in a sufficient amount of detail, relevant to p , so as to rule out the compatibility of the specified details of that world with $\neg p$.
- P3. If proposition p belongs to the class FP (where a proposition is a member of this class iff it is a philosophical proposition that is far-removed from everyday experience), then no-one has imagined a world in a sufficient amount of detail relevant to p .
- C. So, if proposition p belongs to class FP, then one is not justified in asserting the possibility of proposition p .

Recall, from section 3.4, that three objections can be lodged against this argument. The first objection is directed toward P1: why ought we think that the imagination is the only route to justified belief in possibility, particularly Yablo's imaginability? What about plausible suggestions for acquiring modal knowledge, such as through modal intuition, modal logic, the actuality principle and the similarity principle? Recall that I dismissed modal intuition and modal logic as viable competitors to imaginability: the former is simply too subjective, meaning that intractable disagreements over justification loom for the intuitionist; the latter appears to use the c-p principle to form a system of checks and balances against which to test the acceptability of the competing systems of axioms behind various modal logics. Hence modal logic, as a technique, assumes and builds upon imaginability as a guide to possibility. However, the actuality and similarity principle are more difficult to dismiss as reliable means to acquiring modal knowledge.

The second objection is aimed at the proposition "one is justified in asserting the possibility of a proposition p only if one imagines a world in a sufficient amount of detail, relevant to p , so as to rule out the compatibility of the specified details of that world with $\neg p$ " (thus, this objection can be utilized against either P1 or P2 of van Inwagen's argument). The objection is basically that this proposition advocates a theory of modal justification that is far too demanding. There appear to be plenty of uncontroversially true modal propositions that would not pass such a stringent test: the fact that JFK could have avoided the assassins's bullet and gone on to die of a heart attack at a ripe old age was cited as an example, as was the banal fact that my dog could, counter-factually, be lying in the sun right now.

The third objection is aimed against P3. This objection cites the apparent ease with which we can imagine propositions far removed from experience of the actual world, such as those asserting the existence of ghosts, ghouls, goblins and sorcery. What is more, as the fantasy and science-fiction genres attest, we seem to be able to

imagine situations in which such things exist in vast amounts of detail and in such a manner that it is unequivocally true of the imagined world that such things do exist in that world (failure to appreciate this is a failure to understand the facts of the imagined world, not indicative of some kind of philosophical insight that whoever imagined that world somehow got their description of the world wrong).

Let us now re-evaluate these objections in light of the progress we made in the last chapter in formulating a clear and reasonably comprehensive theory of modal epistemology. The results of this re-evaluation will be as follows: the truth of the first and second objection will be granted, leading to a reformulation of van Inwagen's argument that takes on board the insights behind these objections. I will then argue that the third objection is mistaken. In conclusion, a revision of van Inwagen's skepticism will be upheld.

5.1 Can we acquire modal knowledge without Yablo-conceivability?

I am prepared to grant the claim that we can acquire justified modal belief through means other than Yablo-conceivability. In particular, I think we can acquire such justification through the actuality principle and through the similarity principle conjoined with the actuality principle. I do not think it is the case that either of these techniques either i) presuppose Yablo-conceivability or ii) entail, for a proposition justified by one of these techniques, that that proposition is also Yablo-conceivable. I do think that, in general, basic modal claims are easy to imagine (that is, it is easy to imagine a fictional scenario of which they hold true). Nevertheless, imagination need not come into it when the actuality principle supplies justification: I know it is possible that I write this sentence, because I am writing this sentence. No imagination required! The relationship between imaginability and the similarity principle

might be a closer one, however. Indeed, for a great many non-actual possibility-claims, imagining that situation is the means by which we come to think about such a scenario in the first place. Why am I thinking about a burglar entering my house, when it is false that a burglar has entered my house? It cannot be observation or memory that has put this idea in my head. Indeed, it seems that I *made this scenario up*. “Making things up” like this has obvious utility in, for instance, planning and decision-making, because this kind of imagining makes us aware of (presumably) non-actual scenarios. I do not think, however, that this kind of imagining is enough to convince us that the imagined scenario is possible: I think of my burglar scenario as possible because of my *experiences*, because of the similarity principle, not *merely* because I imagined it. It is the similarity principle that does the justificatory work here. The imagination, in this case, simply produces the representation of a non-actual scenario, which then allows us to compare this non-actual situation to an actual one. Further, none of this rules out the possibility that there are *other* means for coming up with representations of non-actual scenarios, or that we always *need* to ponder a non-actual scenario to make a possibility conclusion using the similarity principle (I tend to conclude that drinking glasses are breakable without any obvious use of my imagination). Finally, and perhaps most importantly, the kind of imaginability that is typical of possibility-claims justified by the actuality and similarity principles does *not* seem to be close to Yablo-conceivability: I can imagine that my dog is sleeping in the sun, but I *cannot* construct a fictional scenario in which it is *not* directly stipulated that my dog is sleeping in the sun, but it is rather *implied* by other fictional truths that my dog is sleeping in the sun. This is because, since I am not aware of any of the *essential* (genetic or otherwise) properties of my dog (only certain contingent facts about its history), I cannot construct a scenario that is incompatible with the background scenario that the dog in the imagined scenario is a clone of my dog, that replaced my dog at its birth and has, since, experienced

precisely the same history as my dog in the actual world.

Given that I agree with this objection, where does this leave van Inwagen's skepticism? The main thrust of van Inwagen's argument is not compromised by this objection. Indeed, a simple revision of the argument, one which follows naturally from the "safe explanation" model of modal epistemology, will suffice to shield the argument from the objection: since the conclusion of van Inwagen's argument is directed at *non-basic* modal claims, why not simply rephrase the entire argument in terms of non-basic modal claims? This gels well with the "safe explanation" model, as this theory is committed to something like the following proposition, which is the first premise of the revised skeptical argument:

One is justified in asserting the truth of a non-basic modal proposition "it is possible that p" only if someone has imagined a world the imaginer takes to verify the proposition p.

The objection is impotent against this proposition. Taking "non-basic" to refer, intuitively, to those modal propositions that are far-removed from actual experience, it is far from clear what other means we might have for assessing the truth of these possibility-claims, other than imagination-centered techniques that have been carefully formulated to avoiding the weaknesses of "mere" imagining. Experience-based techniques, such as those formed around the actuality and similarity principle, are (by themselves) obviously inadequate. If any other alternatives exist, then what are they? I certainly can think of none.

5.2 Is Yablo-conceivability too demanding?

With respect to basic possibility-claims, it is granted that Yablo-conceivability is too demanding. It is clearly unrealistic to suggest that we need such potentially complicated imaginative exercises to justify everyday modal claims. Fortunately,

the revision to van Inwagen’s argument that was introduced in the previous section also offers resolution to the charge that Yablo-conceivability is too demanding for basic possibility-claims: if the argument is phrased in terms of non-basic possibility-claims (while in the background, it is acknowledged that basic possibility-claims have a different ground for justification), then worries about basic possibility-claims can safely be set aside¹.

A further problem persists, however: is Yablo-conceivability too demanding for even non-basic possibility-claims? Indeed, as I have spelt out Yablo-conceivability, in the formulation of the “safe explanation” model, I have added further conditions to Yablo-conceivability, on top of the initial claim that Yablo-conceivability involves imagining a world that is incompatible with the falsity of *p*. In other words, I have argued that the criteria for justified belief in non-basic possibility is even stricter than what Yablo and van Inwagen initially suggest. Have I gone too far, eliminating any opportunity for justified belief in non-basic possibility-claims? In particular, is it really feasible to demand a detailed explanation for the truth of a proposition *p* in a fictional world, when the amount of detail required by such an explanation might be significant?

The first thing that needs to be noted is that it isn’t clear that it *ought* to be at all easy to reach the point where one can confidently assert the truth of non-basic possibility-claims. In engineering and the sciences, possibility claims are not taken

¹Does this include the JFK case? This isn’t clear, mainly, I believe, because without additional details, it isn’t clear if the JFK case counts as a basic or a non-basic possibility-claim. If the possibility-claim under consideration is the broad one “it is possible that JFK could have died of a heart-attack”, then this may well qualify as basic modal claim, with justification forthcoming from the similarity principle: as a human being, JFK might well have been relevantly similar to the kinds of people who have died of heart attacks. On the other hand, perhaps JFK was not relevantly similar to people who have suffered from heart-attacks (he was a health and fitness nut and was genetically predisposed to have a strong heart). With this information, it is no longer clear that it is a *basic* claim that he could have had a heart attack - a more detailed story would be required to explain why, counter-factually, JFK could have gotten into a health position that made him relatively prone to heart-attacks. Furthermore, if a more detailed proposition is under scrutiny, such as “it is possible that JFK died of a heart attack in 1989 in Israel, him being on the cusp of brokering a historic peace deal between the Palestinians and Israel”, then it seems clear to me that this is a non-basic modal proposition, which again would require the construction of a detailed fictional world.

lightly, for instance. For the hard-nosed scientist, the possibility of transparent iron is not something we can merely take for granted, following some ill-defined imaginative exercise. Similarly, an engineer is not likely to take seriously the possibility of a working teleporter without non-trivial reasons for backing up such a claim. Both, no doubt, would only be satisfied with substantive explanations as to how such things could be. Why should philosophers and the laymen be exempt from these tough standards?

Furthermore, the “safe explanation” model is perhaps not as strict as it seems. Geirrrson also lodges the complaint against van Inwagen’s skepticism that Yablo-conceivability is too demanding [8]. What he develops from this worry is of some interest: he suggests that Yablo’s conceivability is not fundamentally flawed, but must be amended to take account of various *degrees* of justification. To be precise, he suggests that the more detail, relevant to proposition p , that is added to the imagined scenario, the more evidence one has for asserting the possibility of p . The amount of relevant detail one provides in the imagined world is precisely proportional to the justification one has in asserting the possibility of p . In other words, constructing the imagined world so that it includes *every single* detail relevant to p is not required for justification. There is some merit to this suggestion, and, fortunately, the “safe explanation” model accommodates this insight. The model provides two reasons why more detail in the fictional world is always better: i) the more consistent detail, in general, in the world, the higher the probability that that fictional world is a partial description of a possible world, and, more importantly, ii) the more detailed the explanation for p , the more likely that the offered “explanation” has no gaps or holes in it and indeed acts as an explanation for p . More detail being preferable does not entail that more detail is *required* for justification, however - for practical reasons, we generally have to make do with something less than complete explanation, and, in principle, we must always make do with detail less than that of a possible world.

There thus appears to be some minimum level of detail required for justification, with justification increasing in parallel to the level of detail. I won't attempt to formulate this principle more precisely here. For the purposes of this essay, we can stick to intuitive examples: in order to be convinced that the outrageous misanthropy of a villain such as Hannibal Lector is possible, we require some explanation as to how his character was shaped. An explanation involving vague allusions to childhood trauma goes some way to being persuasive; however, a detailed "origin-story" is even better (witness the recently published novel which chronicles the early life of Lector). As a second example, consider what it would take to be convinced that the Confederates could have won the American Civil War. Substantially convincing would be a broad explanation in terms of certain key battles, upon which the war hinged, being won, instead of lost, by the Confederates. Even more convincing, however, would be an explanation detailed enough to describe the logistics as to *how* the Confederates might have won those battles. In short, it is a normal feature of our explanations that they can serve a justificatory function without being complete - but that, nevertheless, added detail is always more compelling.

5.3 Van Inwagen's argument revised

Hence, the "safe explanation" model appears to put van Inwagen's skepticism in a defensible position. All that is required is a slight revision of the skeptical argument in terms of the reading of Yablo-conceivability encompassed in the model. Here then is van Inwagen's argument revised, in a form impervious to the first and second of the three objections:

van Inwagen's skeptical argument (revision)

- P1. One is justified in asserting the truth of the non-basic possibility-claim "it is possible that p" only if someone has imagined a consistent fictional world such

that propositions p_1, p_2, \dots, p_n hold for that world, where i) p is deduced from p_1, p_2, \dots, p_n by the imaginer and ii) either $\diamond p_k$ ($1 \leq k \leq n$) is a basic possibility-claim, or $\diamond p_k$ is a non-basic possibility-claim that has been justified to some acceptable degree.

- P2. If proposition p belongs to the class FP (where a proposition is a member of this class iff it is a philosophical proposition that is far-removed from everyday experience), then “it is possible that p ” is a non-basic possibility-claim and no-one has imagined a fictional world such that a “modally safe” explanation for p is embedded in the details of that fictional world.
- C. So, if proposition p belongs to class FP, then one is not justified in asserting the possibility of proposition p .

With the advent of a seemingly successful reading of Yablo-conceivability, the first and second premises of the original formulation of the argument have been collapsed into one premise. The only matter that needs to be settled now, before we can safely support the conclusion of this argument, is to defend P2 above.

5.4 Are we able to Yablo-conceive of ghosts, ghouls, goblins and sorcery?

My defense of the second premise of the revised skeptical argument will be an inductive one: I will cast doubt on the claim that anyone has Yablo-conceived of ghosts, or ghouls, or goblins, or sorcery. In other words, I will argue that no philosopher, that I am aware of, has imagined a fairly detailed, consistent fictional world such that embedded in the details of that fictional world is a “modally safe” explanation for any one of the propositions on the following list: “there exists a ghost in this world”, “there exists a ghoul in this world”, “there exists a goblin in this world” and

“the alakazam law holds in this world”. I will accomplish this by giving some rough indication as to what would be required in giving a “modally safe” explanation for such things, and the difficulties that ensue with these requirements.

I will make one (hopefully reasonable) assumption in the forthcoming discussion: I assume that each of the propositions in this list, when placed behind the “it is possible that” operator, forms a non-basic possibility-claim. In other words, none of these propositions can be justified through the actuality or similarity principle (in particular, I assume that the propositions in the list are all non-actual - that is, it is false that ghosts, ghouls, goblins and sorcery exist in the actual world). Yablo-conceivability is thus the only candidate technique for justifying these possibility-claims.

5.4.1 Ghosts

What would it take to give a “modally safe” explanation of the presence of a disembodied mind in a fictional world? We should perhaps first ask ourselves what fictional truths about minds would count as uncontroversial. I will assume that we know from the actual world that minds (whatever they are - and whether they be material or non-material) exist, and, further, that they exist in a close relationship with living human brains. That is, in a fictional world, there is nothing controversial about claiming the existence of a mind, so long as there is a normally functioning human body, with a healthy brain, in which the mind can be “housed”. There is no substantive evidence in the actual world that minds can be separated from living brains - there is no accepted case of a mind that exists housed in some other material object (for instance, a circuit-based computer of some description) and there is no accepted case of a mind existing with no accompanying material object *period* (one might wonder if it is even possible, in principle, to collect evidence in support of the latter).

What remains a tangible mystery is the *precise connection* between living brains and minds. Changes in a living brain co-relate with changes in the consciousness associated with that brain. If certain functions of that brain are damaged, then the mental abilities of the associated mind are impaired in a predictable way. Yet no-one can say *why*. Our inability to understand why minds in actual fact exist (which appears to be a problem explaining how living brains instantiate minds), casts significant doubt on our ability to explain alternative ways in which minds might exist. What is more, if a mind in a fictional world exists under very different circumstances to minds in the actual world, then there is a worry that the explanations that we glean from the actual world might lose their relevance in the fictional case. Our best hope would be to have a comprehensive understanding of minds and an explanation as to how minds are instantiated in the actual world, which, perhaps, might make it clear how the link between minds and material objects could be severed. But we are certainly not privy to this information as it stands.

To illustrate our inability to explain the presence of disembodied minds in a fictional world, consider a fictional world in which a ghost exists (if your imagination fails you, borrow the fictional world from one of the countless Hollywood ghost stories, such as the *The Sixth Sense*). Now, it is required, for this imagining to count as Yablo-conceiving, that all the controversial modal facts in this world have a “modally safe” explanation that emanates from the uncontroversial modal facts. To this end, put the claim that ghosts exist in this world, on hold. Now that it hasn’t been stipulated that ghosts exist in this world, what might lead one to *conclude* that ghosts exist in this world? Remember, being able to conclude, unequivocally, that ghosts exist in this world, implies having a world that is incompatible with the claim that ghosts don’t exist. What, however, in this imagining compels one to assert the existence of a ghost? A supernatural realm cannot be imagined as part of the world - this is just as controversial a claim as the fact that the ghost exists. Ought we

to construct the world so that there are events in that world that are meant to be consequences of the ghost interacting with the physical world? Hollywood movies favour showing objects that move, yet have no physical cause behind their moving. Yet an object that moves with no physical cause is itself a controversial modal claim. Indeed, it is of the same ilk as imagining somebody moving an object with nothing but the power of their mind. Surely, however, few are willing to accept the possibility of telekinesis very easily². In short, I am not convinced that anyone has imagined a fictional world in which ghosts exist according to the conditions required to satisfy Yablo-conceivability.

5.4.2 Ghouls

Considerations similar to the last section cast doubt on our ability to provide an uncontroversial explanation for ghouls in a fictional world. In the actual world some objects (apparently) have minds and others do not. Rocks don't have minds³, clouds don't have minds, and, seemingly, some living human bodies whose brains have been badly traumatized lack minds (witness the phenomenon of the "human vegetable"). On the other hand, normally functioning, living human beings (and maybe animals) have minds. What is the difference between these objects? Although no-one, as far as I know, can explain *precisely* why minds reside in one kind of object and not another, a relevant difference (one that provides the key to our best hope for understanding this phenomenon) is the appearance of a fully functioning human brain as a seemingly exceptionless feature of human beings that manifest the trait of consciousness. Now, again, since no-one can explain exactly how brains instantiate consciousness (and

²Someone might object that every time I flex my hand I am moving an object with nothing but the power of my mind. This point need not affect the argument. There are direct, explicable connections between the muscles in my hand and my brain, and it is given that there is some (mysterious) connection between my brain and my mind. Things only become far-removed from reality when the object moved by the mind, such as doorhandle, has no physical connection to the brain.

³It is perhaps worth noting that panpsychists, such as Leibniz, deny this.

whether they are *uniquely* and *essentially* suited to this role), it seems to me that no one can explain how it might be that a *fully functioning* brain can lack consciousness. How is it that the connection can be severed? How is it that a fully functioning brain can *avoid* instantiating consciousness? We certainly know of cases where living brains reside in bodies for which there is no evidence of an operating mind, but these brains can be unequivocally identified as damaged.

Imagine a situation in which there is a ghoul. The point about this ghoul is that, based on outward appearances, there is no difference between himself and a conscious human being. He displays all the behaviour that we normally use to conclude (unobservable) consciousness. Yet, the fictional fact is that there is no conscious experience for the ghoul: the “lights are out upstairs”. Now, in the interests of ensuring that our fictional world *verifies* the existence of this ghoul, let us put on hold the stipulation that this human being is not conscious. Given this, is there anything about this situation that would entail that this human being is not conscious? Indeed, the fictional facts of our world are more likely to lead one to believe that our ghoul is conscious, since he displays both the behavioral tendencies of conscious being, and any (fictional) medical examination of his brain will reveal a normally working brain. All in all, there is little reason to think this human being is not conscious, let alone verification that he is not conscious. What is required is an explanation as to how, despite his normal physical condition, no mind is instantiated in his body. Such an explanation, I maintain, is not forthcoming.

5.4.3 Goblins

What would it take to imagine a consistent fictional world in which it follows from modally uncontroversial fictional truths that goblins inhabit that world? Let us assume that it would be enough, to show the existence of goblins, if the explanation entails that a species of creature exists in that fictional world that has the following

naturally occurring features: these creatures are somewhat similar to human beings, in build and level of intelligence, but are short (by human standards), green-skinned, pointy-eared, foul tempered and averse to sunlight⁴.

As it turns out, the prospect of providing a modally uncontroversial explanation as to how these creatures exist in this world, is fairly hopeful compared to our previous two cases. It is plain that what we require here is that goblins are a natural biological kind in this world, and, hence, are a species defined by a certain kind of DNA, that guarantees the observable features I listed earlier. Hence, what needs to be explained is how such DNA is possible. Ideally, this requires i) a detailed account as to the structure of this DNA and how this structure leads to the features in question and ii) an account as to how this DNA arose in this world. Rough answers to both questions appear to be forthcoming. With regard to i): goblins are broadly similar to human beings, so it seems that goblin DNA will not be wildly different to a human being's (perhaps they might even count as a sub-branch of the human family?). Furthermore, it is an actual fact that human DNA has enough variety in the genetic information it carries so that human beings have varying heights, skin colour and temperaments. It therefore does not seem a stretch to postulate the hypothesis that a kind of DNA might exist that carries genetic information that ensures the carrier of the DNA has the skin properties and temperament of a goblin. As to ii): one suggestion is to appeal to natural selection in this fictional world: goblins were selected due to their survival traits. Another suggestion might be that they were created through genetic engineering.

This explanation is very broad and rough. Is it enough to meet the minimum

⁴This is a bit like saying that it is sufficient that rhinos exist in a world if there are creatures there that are somewhat elephant-like in build (except stockier), have a grey, thick hide and a horn protruding vertically from the end of their nose. It seems to me that having creatures that fit this description is not enough to guarantee that there are rhinos in that fictional world. Does a list of typical-for-the-species, naturally occurring features suffice for a non-actual creature like a goblin or a unicorn, however? Perhaps not. Indeed, Saul Kripke, for instance, thinks that purely fictional creatures are necessarily non-existent [10, pp.156-158]. None of this need concern us here, however.

level of explanation needed to justify belief in the possibility of goblins? I don't think so. What has been said certainly indicates the direction one would need to go in, if one were serious about confirming the possibility of goblins (and who would be?). It seems to me that a great deal more would need to be said to really explore this possibility. Is human-like DNA possible that can produce a creature with skin that is the deep shade of green fitting for a goblin? Can this DNA also produce skin that is sensitive to the light? Can it also produce ears that are as hideously long and pointy as the goblins? Is it really the case that, if genetic engineering were to become a reality, one would have the option to engineer, perhaps for a laugh, one's offspring to become a goblin?

The explanation in connection to ii) is also hopelessly shallow at this stage and it isn't clear to me how it ought to be filled out (I am also fairly certain that no-one has attempted such a feat in any seriousness). Why would natural selection have produced goblins in this world? Providing an answer would require a reasonably detailed account of the kind of environments present in this world. Perhaps establishing the motivation of a mad genetic engineer would be easier - but this route requires some explanation as to how genetic engineering is possible in the first place.

In short, I think we have some idea as to how to provide an adequate explanation as to how goblins exist in a fictional world. I do, however, doubt that anyone has pursued this idea to fruition. One may be wondering at this point why we are discussing goblins at all, of course. My purpose is an inductive one: if we haven't provided an explanation for the existence of non-actual creatures such as goblins, creatures we seem to be able to directly imagine in a fair amount of detail, why think it will be any easier to provide one for the existence of even more ambitious creatures, such as those with infinite minds, those with perfect rationality, utility monsters, God or any of the other members of the cast of (presumably) non-actual beings the philosopher likes to wheel out?

5.4.4 Sorcery

The “Alakazam” law is as follows: if one holds some wormwood in one’s left hand, points at an object with one’s right index finger and utters the word “alakazam!”, then whatever is in the line of sight of one’s finger will promptly burst into flames. What kind of explanation can one offer for a world in which the “Alakazam” law holds?

An advocate of Humean supervenience might have a ready answer. Humean supervenience is the doctrine that the laws of nature are supervenient on the history of the world - the laws of nature simply organize or summarize the regularities present in the particular matters of fact that make up that history. In this case, imagining a fictional world in which it follows from the fictional truths of that world that the laws of nature are different is easy: simply imagine a world that has a different history to the actual world, a history that supports different laws of nature [20, pp.317-318].

The problem with this suggestion, however, is that this explanation can hardly be called “modally safe”. The particular matters of fact that are meant to support the non-actual laws of nature are themselves hardly modally uncontroversial. It appears equally controversial to stipulate of a fictional world that the particular (causal) event “Sue held some wormwood in her left hand, pointed at an object with her right index finger and uttered the word “alakazam!”, causing something in the line of sight of her finger to promptly burst into flames” occurred there, as to stipulate the general Alakazam law.

So an appeal to Humean supervenience is of little help. We need to justify the existence of a law without appeal to particular instantiations of that law. However, I am doubtful that anyone is able to offer such an explanation for a law such as the “Alakazam” law. It is simply too arbitrary. If arbitrary laws of this ilk were suitably justifiable and, hence, modally uncontroversial, then it would follow that *anything* could be explained in a “modally safe” fashion in a fictional world - simply set up the

world so that it is governed by an appropriate law (if one wants to explain the truth of p in that world, simply set up a law in that world so that some modally uncontroversial event is guaranteed to be followed by an event to the effect that p becomes true in that world). However, it is clearly absurd to think that such a procedure has any effective explanatory power. To illustrate this, consider the following three laws:

Alakazam Law 1: Whenever the word “Alakazam” is uttered, a ghost (similar to Casper the friendly ghost) pops into being.

Alakazam Law 2: Whenever the word “Alakazam” is uttered, the utterer of the word is immediately transformed into a ghoul/zombie.

Alakazam Law 3: Whenever the word “Alakazam” is uttered, a goblin pops into being.

I can imagine a fictional world in which these laws hold just as easily as I can imagine a Humean fictional world in which the mechanics of billiard tables are otherwise to those of the actual world. Furthermore, these three laws can act as the basis for an explanation for the existence of ghosts, ghouls or goblins in the respective worlds in which they hold. Does this then solve the problem of providing a “modally safe” explanation of the existence of ghost, ghouls and goblins? Clearly not. The discussion in the last three sections ought to make one doubtful that such easy answers are forthcoming. In short, it is clear that, for exceedingly arbitrary choices of laws in fictional worlds, these laws holding is modally controversial, and, what is more, the arbitrariness of these laws leaves little hope for an explanation for their instantiation.

The above indicates, at the very least, that if we are justified in believing that the laws of nature could have been different, then we are certainly *not* justified in thinking that just any old laws could have held. However, perhaps more conservative suggestions for laws, alternative to those in actual world, provide options for explanation that laws of the Alakazam ilk do not. Let us consider.

The only way to explain a given law of nature, it seems to me, is to make an appeal to even more general laws of nature. This appears to be the aim of physics - to ultimately reduce the state of the physical world to interactions between fundamental elements governed by fundamental laws. Now, the advocate for contingent laws of nature might offer the following argument. Given any particular law of nature, an explanation in terms of “higher-level” laws will suffice. At some point, however, the most fundamental laws of that world need to be stipulated. No explanation, however, can be asked for these fundamental laws. Indeed, since these fundamental laws are meant to, in some sense, explain *everything* physical about that world, there is little sense in asking for an explanation *for them*. They represent where the explanations stop. What is more (the advocate of contingent laws continues), it is safe to say that the fundamental laws in a fictional world are modally uncontroversial, even if those fundamental laws are somewhat different to the actual fundamental laws. Why? Well, when we discover the fundamental laws and fundamental elements that underly the actual world (let us say, for instance, that we discover five fundamental particles that interact according to 120 fundamental laws), then what we have discovered is a way the actual world simply *is*. No further explanations for why the world is like that are forthcoming. Hence, there is no *reason* why the actual world is like that. However, since there is no reason why the world is like that, there is no reason why the world *could not* have been like that. Hence, it is modally uncontroversial that the world could have been governed by very different fundamental particles, interacting according to other laws.

Now, assuming that the above line of reasoning is correct, and we thus have a rough strategy indicating how we can explain changes in the laws of nature, it becomes apparent that, according to the “safe explanation” model, in order to justify a “low-level” counter-factual law such as that “metals shrink when heated” or “signals can travel faster than the speed of light, but not faster than twice the speed of

light”, what is required is an explanation for that law in terms of a picture of the fundamental elements and interactions of that world. However, working out the observable consequences of such a fundamental picture is difficult enough - working backwards from a “low-level” law that one wants to justify to “higher-level” picture that does indeed justify that law is even more difficult. In other words, explaining a different set of laws in a fictional world ought not to be taken lightly. These considerations are enough, I think, to cast doubt on the idea that anyone has imagined a world in which “metals shrink when heated” is true, or where the laws of mechanics governing billiard balls are otherwise to the actual world, let alone one where the “Alakazam” law holds.

With this, my inductive defense of the second premise of the revised skeptical argument is complete. Hence, I advocate the truth of the conclusion of the revised skeptical argument.

Chapter 6

Conclusion

6.1 Summary of results

In this paper, I defended the following skeptical argument, inspired by an argument from van Inwagen:

van Inwagen's skeptical argument (revision)

- P1. One is justified in asserting the truth of the non-basic possibility-claim “it is possible that p ” only if someone has imagined a consistent fictional world such that propositions p_1, p_2, \dots, p_n hold for that world, where i) p is deduced from p_1, p_2, \dots, p_n by the imaginer and ii) either $\diamond p_k$ ($1 \leq k \leq n$) is a basic possibility-claim, or $\diamond p_k$ is a non-basic possibility-claim that has been justified to some acceptable degree.
- P2. If proposition p belongs to the class FP (where a proposition is a member of this class iff it is a philosophical proposition that is far-removed from everyday experience), then “it is possible that p ” is a non-basic possibility-claim and no one has imagined a fictional world such that a “modally safe” explanation for p is embedded in the details of that fictional world.

- C. So, if proposition p belongs to class FP, then one is not justified in asserting the possibility of proposition p .

I defended the first premise by defending the following recursive account of modal epistemology, what I call the “safe explanation” model:

A basic possibility-claim is justified iff it is supported by either the actuality principle, or the similarity principle in conjunction with the actuality principle.

A non-basic possibility-claim “it is possible that p ” is justified iff someone has imagined a consistent fictional world such that propositions p_1, p_2, \dots, p_n hold for that world, where i) p is deduced from p_1, p_2, \dots, p_n by the imaginer and ii) either $\diamond p_k$ ($1 \leq k \leq n$) is a basic possibility-claim, or $\diamond p_k$ is a non-basic possibility-claim that has been justified to some acceptable degree.

This account is an expansion of the suggestion made by Yablo as to the kind of conceivability that is relevant to guiding our possibility-claims:

p is [philosophically] conceivable for me iff [f] I can imagine a world that I take to verify p [29, p.29].

The account of how basic possibility-claims are justified, according to the “safe explanation” model, rests upon the triviality of the actuality principle and a defense of the similarity principle in terms of the principle of induction and the actuality principle. The account as to how non-basic possibility claims are justified, rests on the grounds that imagining a fictional world according to the restrictions outlined in the account will guarantee the highest probability that that fictional world partially describes a possible world, since the fictional world will be i) as consistent as a fictional

world can be, ii) reasonably detailed and iii) constructed from uncontroversial modal claims.

I defended the second premise inductively, by working through a list of representative possibility-claims (namely, the possibility that ghosts exist, the possibility that ghouls exist, the possibility that goblins exist and the possibility that sorcery exists) and showing that none of these possibility-claims are justified, as far as I can see, according to the “safe explanation” model.

6.2 Avenues for further research

I make the following proposals for further research:

- i. One might wonder where these results leave some of philosophy’s most important and influential thought experiments, since the fictional worlds that ground these experiments rarely meet the requirements laid down to justify them as representative of possible worlds. An avenue that I think is worth exploring, in this regard, is that many of these thought experiments need *not* be construed as discussing possible worlds at all. The possibility of the features of the fictional world might be irrelevant, since the aim of such experiments appears to be a semantic or conceptual one: to hear what one would say about such a world, whether it be possible or not. Hence, I think it is worth exploring the idea that philosophical thought experiments, unlike the thought experiments of physics, are mainly aimed at drawing conclusions about our language and concepts, not necessarily metaphysical conclusions about the state of reality. This view would fit well with the results of this paper and save many philosophical thought experiments from the rubbish heap.
- ii. One thing that is lacking from my account of the “safe explanation” model is a really rigorous account as to how to assign a degree of justification to a

- possibility-claim. I am hopeful of the possibility of such an account: let us say that the degree of justification for a possibility-claim is between 0 and 1, where 0 means that there is no justification for it, while 1 means that it is confirmed that that possibility-claim is true. Now candidates for being assigned a 1, might be some of the claims that are justified by the actuality principle, or perhaps some non-basic claims that have been justified according to very comprehensive and rigorous explanation (such as by producing a very detailed blueprint). How though to compute degrees of justification in general? The justification of the propositions which explain a non-basic possibility-claim is meant to be “transferred” to the non-basic claim. Can this “transference” be formalized?
- iii. It might be worth exploring further the first option I presented, in section 4.2.2, for accounting for basic modal claims - namely, the idea that “branching points” are created in the history of the world by the free choices of free agents. Exploring this idea takes one deep into the debate concerning the nature of freedom, so is best left as a separate research project.

Bibliography

- [1] CHALMERS, D., *The Conscious Mind*, Oxford University Press, 1996.
- [2] CHALMERS, D., Does Conceivability Entail Possibility?, in *Conceivability and Possibility*, pp. 145-200, Hawthorne, J., Gendler, T.S. eds., Oxford University Press, 2002.
- [3] DESCARTES, R., *Meditations*.
- [4] ELLIS, B., *The Philosophy of Nature: A Guide to the New Essentialism*, Queens-McGill Press, 2002.
- [5] ESCHER, M.C., p.69, pp.74-76, *The Graphic Work of M.C. Escher*, Hawthorn Books, London, 1960.
- [6] FINE, K., “The Varieties of Necessity”, in *Conceivability and Possibility*, pp. 253-282, Hawthorne, J., Gendler, T.S. eds., Oxford University Press, 2002.
- [7] GEDDON, G., “Logical Possibility”, *Mind*, New Series, Vol.81, No.324, pp.481-494, October 1972.
- [8] GEIRSSON, H., “Conceivability and Defeasible Modal Justification”, *Philosophical Studies* 122, pp.279-304, 2005.
- [9] HUME, D., *A Treatise on Human Nature*, p.32, Selby-Bigge, L.A. ed., Oxford University Press, London, 1965.

- [10] KRIPKE, S., *Naming and Necessity*, Harvard University Press, Cambridge: MA, 1980.
- [11] LEMMON, E. J., "Is There Only One Correct System of Modal Logic?", *Aristotelian Society*, Supplementary Volume, pp.23-40, Vol.33, 1959.
- [12] MACKIE, P., *How Things Might Have Been: Individuals, Kinds, and Essential Properties*, Oxford University Press, 2006.
- [13] MILL, J.S., *A System of Logic*, in *Collected Works of John Stuart Mill*, VII, p. 269, Robson, J.M. ed.
- [14] PRIEST, G., *In Contradiction*, Martinus Nijhoff, 1987.
- [15] PUTNAM, H., "The Meaning of 'Meaning'", in his *Mind, Language and Reality*, p. 233, Cambridge University Press, Cambridge: England, 1975.
- [16] QUINN, P.L., "Metaphysical Necessity and Modal Logics", *The Monist*, Vol. 65.
- [17] REID, T., *Essays on the Intellectual Powers of Man*, pp.431-432, MIT Press, Cambridge: MA, 1969.
- [18] SALMON, N., "The Logic of What Might Have Been", pp.3-8, *The Philosophical Review*, Vol. 98, No.1, Jan. 1989.
- [19] SHOEMAKER, S., Causal and Metaphysical Necessity, *Philosophical Quarterly*, Vol 79, No. 1, pp. 59-77, 1998.
- [20] SIDELLE, A., "On the Metaphysical Contingency of the Laws of Nature", in *Conceivability and Possibility*, pp.309-336, Gendler, T.S., Hawthorne, J. eds., Oxford University Press, 2002.
- [21] SORENSEN, R., "The Art of the Impossible", in *Conceivability and Possibility*, pp.337-368, Gendler, T.S., Hawthorne, J. eds., Oxford University Press, 2002.

- [22] SORENSEN, R., "Meta-conceivability and Thought Experiments", in *The Architecture of the Imagination*, pp. 257-272, Nichols, S. ed., Oxford University Press, 2006.
- [23] SORENSEN, R., *Thought Experiments*, Oxford University Press, 1992.
- [24] SOSA, E., "Modal and Other A Priori Epistemology: How Can We Know What is Possible and What is Impossible?", *The Southern Journal of Philosophy*, Vol. XXXVIII, Supplement, 2000.
- [25] STALNAKER, R., "What is it Like to be a Zombie?", in *Conceivability and Possibility*, pp.385-400, Gendler, T.S., Hawthorne, J. eds., Oxford University Press, 2002.
- [26] TIDMAN, P., "Conceivability as a Test for Possibility", *American Philosophical Quarterly*, Vol. 31, No. 4, pp.297-309, October 1994.
- [27] VAN INWAGEN, P., "Modal Epistemology", *Philosophical Studies* 92, pp.67-84, 1998.
- [28] WALTON, K., *Mimesis as Make-Believe*, pp.11-67, Harvard University Press, 1990.
- [29] YABLO, S.C., "Is Conceivability a Guide to Possibility?", *Philosophical and Phenomenological Research*, Vol. LIII, No.1, pp.1-42, March 1993.
- [30] YABLO, S.C., "Coulda, Woulda, Shoulda", in *Conceivability and Possibility*, pp.441-492, Gendler, T.S., Hawthorne, J. eds., Oxford University Press, 2002.