
TWO PROBLEMS IN DYNAMIC ETHICS

COURTNEY M. COX

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY AT THE UNIVERSITY OF OXFORD



CORPUS CHRISTI COLLEGE
UNIVERSITY OF OXFORD
TRINITY TERM 2011

ABSTRACT

TWO PROBLEMS IN DYNAMIC ETHICS

COURTNEY M. COX
CORPUS CHRISTI COLLEGE

THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY,
UNIVERSITY OF OXFORD, TRINITY TERM 2011

Time raises a host of difficult ethical questions. This doctoral project focuses on two:

1. How are “static” comparative principles (e.g. equality, desert) to be understood over time? (The Problem of Fairness & Time)
2. How might separation (in time) between agents, objects, and threats affect claims to the relevant resources? (The New Problem of Temporal Distance)

My work begins with a simple observation: our *prima facie* intuitions about the value of simple distributions change depending on whether such cases are presented as static (occurring at one time) or dynamic (extended over time). Further examination of more complicated distributions leads to the proposal of a new theory, Weighted Progressive Egalitarianism. This theory has two features: only past-regarding complaints matter (a scope restriction), and a comparative complaint between persons located at a great temporal distance matters less than a complaint between contemporaries (a weighting restriction).

This theory provides one plausible answer to the first question, the Problem of Fairness & Time. The evaluation of this theory relies on and reveals some non-standard answers to the second question, the New Problem of Temporal Distance. I conclude by arguing that the theory’s application to a few puzzles in population axiology merits further investigation.

DECLARATION

All the work in this thesis is my own. Part I is developed from my BPhil Thesis. The discussion of the argument from overlapping generations has also been revised from the earlier thesis. In addition, a small portion of the discussion about desert claims to well-being has been developed from a BPhil examination essay.

There were approximately 73372 words prior to the completion of minor corrections.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the intellectual and moral support of a number of friends, family, and colleagues. It would also not have been possible without the financial support of the Oxford University Press's Clarendon Fund Scholarship, and Corpus Christi College's A E Haigh Graduate Scholarship.

There are few people who deserve special thanks. First, I would not have written this thesis had I not been encouraged to attend Oxford. For that, I owe special thanks to Shelly Kagan and Samuel Gorovitz, with whom I first began working on these ideas.

Second, many of my fellow graduate students, too many to name individually, were immensely helpful in discussing the problems and questions contained herein. Special thanks are owed to Marco Grix, who offered a host of very helpful comments on an earlier draft of Part I. I also owe a tremendous debt to Olivia Bailey and Rob Watt for many useful conversations and emails exchanges about the material in Parts II and III.

Third, I thank the reviewers of the BPhil thesis, which was the basis for Part I of this work; my confirmation of status reviewer, for their generous and helpful comments; and my DPhil examiners, Ralph Wedgwood and Gustaf Arrhenius, for useful leads and detailed suggestions.

Above all, I am especially indebted to my supervisor, John Broome. His comments on both the current draft of this work, and many earlier versions, have been critical, insightful, and generous in number. He has been exceptionally patient and generous with his time. He has saved me from making many mistakes; any that remain are my own.

Lastly, this thesis would not have been possible without the moral support of friends and family. I owe special thanks to Murray Tipping, for a steady supply of tea and chocolate; to Thea Goodsell, who packed my freezer with ready-made meals to get me through the final stages; to my grandparents and brother, for their unwavering support; and to my parents. My mother, Jean Cox, has helped to keep me focused with encouragement and kind words. My father, Bruce Cox, read through several drafts of this work, despite finding equality "very complicated". He remains my first and most trusted editor.

CONTENTS

1	INTRODUCTION	1
I	The Problem of Fairness & Time, Part I	8
2	PRELIMINARIES	9
2.1	PRELIMINARIES I: WELL-BEING	9
2.2	PRELIMINARIES II: THE GOOD AND THE MATH	14
2.3	PRELIMINARIES III: TEMPORAL UNITS OF COMPARISON	15
3	THE SIMPLE TEMPORAL CASE	16
3.1	SPURIOUS SOURCE I: GROWING POPULATIONS	17
3.2	SPURIOUS SOURCE II: CAUSAL ASYMMETRY	19
3.3	SPURIOUS SOURCE III: EPISTEMIC CONSIDERATIONS	23
3.4	SPURIOUS SOURCE IV: GROWTH OF RESOURCES & TECHNOLOG- ICAL PROGRESS	24
3.5	SPURIOUS SOURCE V: PSYCHOLOGICAL PREFERENCE FOR PROGRESS	25
4	EQUALITY	27
4.1	TYPES OF EGALITARIANISM	27
4.2	COMPLAINTS, AND MEASURING THE DISVALUE OF INEQUALITY . .	34
4.3	EQUALITY, THE SIMPLE CASE, & THE SIMPLE TEMPORAL CASE .	37
4.3.1	ADJUSTING $v_{i,j}$: WEIGHTING FUTURE WELL-BEING	41
4.3.2	ADJUSTING V : ACCOMMODATING A PREFERENCE FOR PROGRESS?	43
4.3.3	ADJUSTING $c_{i,j}$: QUESTIONING THE SCOPE OF INDIVIDUAL COMPLAINTS	45
4.4	THE COMPLEX TEMPORAL CASES & THE PIGOU-DALTON CON- DITION	53
4.4.1	AN OBJECTION: NO DIFFERENCE BETWEEN DIACHRONIC INEQUALITIES AND GENERATIONAL INEQUALITIES	62
4.4.2	THE COMPLEX TEMPORAL CASE: IMPLICATIONS FOR PE & TNE	66

II	The Problem of Temporal Distance	80
5	AN OUTLINE OF A TOY THEORY OF CLAIMS	84
5.1	WHAT IS A CLAIM?	84
5.1.1	WHY NOT RIGHTS?	87
5.1.2	WHY NOT ENTITLEMENTS?	88
5.2	TYPES OF CLAIMS: AN OVERVIEW	88
5.3	DESERT CLAIMS	89
5.3.1	MORAL WORTH CLAIMS: DESERT CLAIMS TO WELL-BEING	92
5.3.2	DERIVATIVE DESERT CLAIMS TO RESOURCES: CLAIMS OF NEED	101
5.3.3	DIRECT DESERT CLAIMS TO RESOURCES: HARVEST CLAIMS	107
5.4	CLAIMS OF CONVENTION	108
5.5	PROXIMITY CLAIMS	113
6	AGAINST TEMPORAL PROXIMITY IRRELEVANCE	114
6.1	CLAIMS ARE TO THE ORIGINAL MUSHROOM	116
6.1.1	MUSHROOM EXCESS: A NEW RESOURCE	118
6.1.2	MUSHROOM EXCESS: STILL THE ORIGINAL RESOURCE, JUST BIGGER	125
6.2	CLAIMS ARE TO THE MUSHROOM EXCESS GROWN IN A GIVEN YEAR	126
6.3	CLAIMS ARE TO A SHARE OF THE HISTORIC MUSHROOM AGGRE- GATE	128
6.3.1	INTERFERENCE FROM MALICIOUS AGENT	133
6.3.2	NO CLAIMS	137
7	TEMPORAL PROXIMITY RELEVANCE: THE NEW PROBLEM OF DISCOUNT- ING IN MORALITY	140
7.1	PRELIMINARIES	141
7.1.1	WELL-BEING & TIME	141
7.1.2	METAPHYSICAL ASSUMPTIONS	142
7.1.3	TIME-INDEXING CLAIMS	143
7.2	TIME AND EXISTENCE OF THE OBJECT NEEDED	146
7.2.1	MULTIPLE AVAILABLE ORGANS	151
7.2.2	CLAIMS TO PARTICULAR OBJECTS OR TO DESCRIPTIVE (AN OF-THE-SORT CLAIM)	155
7.2.3	EXISTENCE OF THE OBJECT NEEDED	158
7.3	TIME AND EXISTENCE OF THE CLAIMANT	163
7.3.1	YES: CLAIMS TO RESOURCES AT TIMES <i>Before</i> THE TIME OF EXISTENCE OF THE CLAIMANT	164
7.3.2	NO: NO CLAIMS TO RESOURCES AT TIMES <i>Before</i> THE TIME OF EXISTENCE OF THE CLAIMANT	168
7.3.3	DERIVATIVE DESERT CLAIMS TO RESOURCES AFTER TIME OF EXISTENCE OF CLAIMANT	170
7.4	CONCLUSION	172

III	The Problem of Fairness & Time, Part II	173
8	FAIRNESS, CLAIMS, AND THE DISVALUE OF INEQUALITY	175
8.1	CLAIMS & FAIRNESS	176
8.2	CLAIMS, FAIRNESS, & COMPLAINT BASES	178
8.2.1	STRENGTH OF CLAIMS AND SIZE OF COMPLAINTS	184
9	FAIRNESS & TIME	187
9.1	THE ARGUMENT	187
9.1.1	THE ARGUMENT AGAINST COMPLAINTS RELATIVE TO THE FUTURE	189
9.1.2	THE ARGUMENT: YES, CLAIMS CAN BACKTRACK PRIOR TO CLAIMANT'S EXISTENCE	190
9.1.3	THE ARGUMENT: NO, CLAIMS CANNOT BACKTRACK PRIOR TO CLAIMANT'S EXISTENCE	192
9.1.4	THE ARGUMENT: CLAIMS WEAKENED BY BACKTRACKING	195
9.2	OBJECTIONS	197
9.2.1	IMPARTIALITY & RELATIVITY: AN OBJECTION TO WEIGHT- ING COMPLAINTS?	197
9.2.2	IS WEIGHTED PROGRESSIVE EGALITARIANISM REALLY AN EGALITARIAN VIEW?	198
9.2.3	TCT IS NOT ABOUT TIME	200
10	FORMULATING WEIGHTED PROGRESSIVE EGALITARIANISM	204
10.1	COMPLAINTS	205
10.2	THE BEST-OFF PERSON (BOP) & THE LESSON OF G-TYPE HIS- TORIES	208
10.3	"LOCALIZING" THE BEST-OFF PERSON (BOP)	214
10.4	A DIFFERENT WAY TO MODIFY BOP: TEMPORALLY SELECTIVE BOP	220
10.5	TEMPORALLY MODIFIED ALL THOSE BETTER-OFF (ATBO) . . .	222
10.5.1	POTENTIAL PROBLEMS WITH TEMPORALLY MODIFIED ATBO	224
10.6	CONCLUSIONS: LESSONS LEARNED MOVING FORWARD	228
11	WEIGHTED PROGRESSIVE EGALITARIANISM: IMPLICATIONS	230
11.1	THE ARGUMENT FROM OVERLAPPING GENERATIONS	230
11.2	WPE, THE NEUTRAL LEVEL & CONTEXT SENSITIVITY	234
11.2.1	THE NEUTRAL LEVEL	234
11.2.2	AN ARGUMENT FOR A CONSTANT NEUTRAL LEVEL	235
11.3	COMPARATIVE COMPLAINTS & THE NEUTRAL LEVEL	242
11.4	WILL THE FLOATING NEUTRAL LEVEL MAKE A DIFFERENCE? . .	246
11.5	AN OBJECTION & A SOLUTION	249
11.6	CLOSING REMARKS	253
	BIBLIOGRAPHY	255

IV	Appendices	262
A	APPENDIX - MUSHROOM GROWTH MODEL	263
A.1	SIMPLE MUSHROOM GROWTH MODEL (CONTINUOUS TIME)	263
A.1.1	SIMPLE MUSHROOM GROWTH MODEL: WITHOUT INTER- FERENCE	263
A.1.2	SIMPLE MUSHROOM GROWTH MODEL: WITH CONSUMPTION	264
A.1.3	INDIVIDUAL LEVELS OF CONSUMPTION	265
A.1.4	CLAIMS TO THE AGGREGATE	266

LIST OF FIGURES

1.1	The Simple Case	1
1.2	The Simple Temporal Case	3
2.1	Road from Oxford	12
3.1	Indefinite End of Humanity	18
3.2	Temporal Leveling Down	22
4.1	Temporal Leveling Down Objection Revisited	45
4.2	Static Leveling Down Objection	46
4.3	Sine-step	49
4.4	Cosine-step	50
4.5	Level	50
4.6	Up-step	51
4.7	(g^*)	56
4.8	(h^*)	56
4.9	(i^*)	56
4.10	Increasing Progression	57
4.11	Repeating Progression	57
4.12	Inheritance Case: Original Distribution	60
4.13	Inheritance Case: (A)	60
4.14	Inheritance Case: (B)	61
4.15	(A')	72
4.16	(B')	73
4.17	(A'')	74
4.18	(B'')	75
10.1	History F	210
10.2	History G	211
10.3	History H-1	216
10.4	History H-2	217
10.5	History H-3	219
10.6	History I	220
10.7	Static Cases J & K	224
10.8	History J	226

10.9 History K	227
11.1 Overlapping Generations	231
11.2 Populations A, B, and C	248

LIST OF TABLES

4.1	TNE and PE Evaluations of Cases	67
4.1		68
5.1	Types of Claims	90

LIST OF ABBREVIATIONS

AD	Absolute Desert
AP	Additive Principle of Equality
ATBO	(complaints relative to) All Those Better Off
AVE	(complaints relative to) the Average
BOP	(complaints relative to) the Best-off Person
GE	Generational Equality
MP	Maximin Principle of Equality
PE	Progressive Egalitarianism
PSR	Progressive Scope Restriction
TCT	Temporal Comparison Thesis
TNE	Temporally Neutral Egalitarianism
TPI	Temporal Proximity Irrelevance
TVV*	Temporally Variable Noncomparative Value of Well-being (a type of temporal bias)
WAP	Weighted Additive Principle of Equality
WPE	Weighted Progressive Egalitarianism

INTRODUCTION

Consider three distributions of welfare:




Figure 1.1: The Simple Case

Individuals are represented as lined up along the x -axis. All the individuals are the same with regard to any morally relevant respect except location and identity:

all have the same moral status, rational capacities, desert, etc. Their welfare, or well-being, is represented vertically along the *y-axis*.¹ Each of (a), (b) and (c) shows a possible distribution of a given amount of well-being over the same number of people.² Thus, the total amount of well-being is the same in (a), (b), and (c), and so too is the average level of well-being.³ Which distribution is best?

Prima facie, the answer is (b). But why? The answer appears to be because (b) is more equal than either (a) or (c). That is, from the perspective of Equality, (b) is better than (a) and also better than (c), and, everything else being equal, Equality – of some sort, or something like it – matters.⁴

How do (a) and (c) compare? From the perspective of Equality, one is not better than the other: both contain the same inequality. (a) and (c) differ only in one of two possible respects: the same people have the same well-being, but are arranged in a different order from left to right; or if the x-axis corresponds to the identity of people, the only difference is who gets more and who gets less. Assuming no morally relevant differences between individuals, the same inequality is present in both cases. And this seems to fit with *prima facie* intuitions about the distributions: as long as we can assume impartiality about the people involved, we too are indifferent between (a) and (c). If morality is impartial between people, and if there are no morally relevant differences between individuals, then there is no moral difference

¹I use the terms “welfare” and “well-being” interchangeably. What constitutes well-being is a matter for another discussion. Distributive principles might be more appropriately applied to some other type of quantity (e.g. resources, opportunity).

² To avoid confusion, assume also that the people are the same people, although they need not be since the point of interest at the moment is the *distribution* of well-being.

³ I have equalized the number of people and total amount of welfare to avoid confusing intuitions. The same question could be asked of distributions of differing sizes. The focus of this discussion is on the distributions themselves (their “shape”), and not all-things-considered judgments that might be affected by the total amount of well-being or total number of people.

⁴The verdict that (b) is better than both (a) and (c) seems true on most any egalitarian principle, or application of another distributive principle in which the features of the individuals (e.g. desert, or in the case of distribution of resources, need) are all equal. A possible exception is a social minimum theory. I set this complication aside. It seems, intuitively, that the satisfaction of at least some distributive principle has at least some intrinsic worth and provides, at the very least, a *pro tanto* reason to prefer (b) over (a) or (c).

between (a) and (c).

But what if, instead of representing the distribution of well-being between people arranged from left to right, the graphs represented the distribution of well-being between people *over* time arranged from earlier to later? What if (a) represents a series of people where each generation is successively better off, (b) where there is no change, and (c) where each generation is successively worse off? Assume, for simplicity, non-overlapping lifetimes, and that the distribution of people is even across time – the size of a generation (the population at a time) remains constant – and that the total historic population is fixed (history ends where the graphs do). Assume also that the only difference between generations is the time at which they live, some earlier and some later. They are all equally deserving, hard-working, virtuous, etc. – that is, morally equal in all respects (except time of existence). Call this the *Simple Temporal Case*.




Figure 1.2: The Simple Temporal Case

It is no longer clear that (b^*) is the best distribution. While equality between contemporaries is appealing, equality between modern society and cave people does not intuitively matter – and not simply because nothing can be done about it. Nothing can be done about many instances of modern inequality between contemporaries, but such inequalities are still intuitively regrettable in a way that the inequality between a caveman and the reader is not. Similarly, an inequality between individuals in the distant future who are better off and individuals now seems less troubling than inequality between contemporaries: why should an inequality stemming from an improvement in the human condition matter? It is no longer obvious that (b^*) is the best distribution because the inequality in (a^*) is not as *prima facie* objectionable as the inequality in (a) . Indeed, some might hold the *prima facie* intuition that (a^*) is at least as good as, if not better than, (b^*) : inequality created by progress over time does not matter.⁵

More enlightening, however, is that the indifference between (a) and (c) in the Simple Case is counterintuitive in the temporal case between (a^*) and (c^*) . In (a^*) , each generation becomes better off: quality of life improves over time. But in (c^*) , each generation is successively worse off. *Prima facie*, the inequality in (c^*) is worse than the inequality in (a^*) . This is an intuition which is perhaps less widely shared, though it is more striking for those who do hold it since the only difference between the two inequalities is arrangement in time.⁶

These intuitions are, at the least, not egalitarian, and, if (a^*) is preferable to

⁵ This intuition appears to be a weaker version of the one underlying the Egyptology Conclusion or the Independence of the Utilities of the Dead axiom: roughly, the intuition that the contributive value of an individual's life cannot depend on facts about the existence or well-being of those who lived in the distant past. *cf.* Arrhenius and Bykvist (1995, 110ff.); Blackorby et al. (2005, Ch. 8); Parfit (1984, 420ff.).

⁶ Something like this intuition might be found in responding to Mere Addition: there is an intuitive resistance to adding individuals to a population who are worse off than that population (perhaps on the grounds of inequality), but not to adding individuals to a population who are better off than that population.

(b*), are anti-egalitarian. As regards Equality, (b*) is better than (a*) and (c*), and there is no difference between (a*) and (c*). But does the satisfaction of Equality matter in the temporal case? In terms of Equality, (b*) may be better than (a*) and (c*) – but does this provide a (*pro tanto*) reason to prefer (b*)? That is, is there any intrinsic⁷ worth to an equal distribution of well-being *over time*?

This question is one version of the issue at the heart of this thesis, what I call the *Problem of Fairness Over Time*: how can (or should) static “comparative” or “distributive” moral principles be extended over time? On an axiological reading, these principles are really about how well-being – in both its quantity and distribution – contributes to the betterness of a state of affairs. The betterness ranking of a distribution considered as a whole is thus taken to be determined by the amount of well-being and how it is arranged over people. Different views have different merits, and the question of which works best is far from settled. Common to them all is the problem of extension over time. As I will discuss, what works at-a-time may face considerable difficulties over time – and not just for reasons related to the now-infamous Nonidentity Problem.⁸

Before continuing, a few clarifications may be useful. There are a number of instrumental considerations that could be driving the anti-egalitarian or indifferent-egalitarian intuitions in the temporal case sketched above. Underlying the preference for (a*) might be the belief that populations grow over time, so that greater well-being for later generations means a greater total amount of well-being. Alternatively, intuitions might be affected by the idea that sacrifices on the part of earlier

⁷ Throughout I use “intrinsic value” and “intrinsic worth” to mean *valuable-in-itself*. This is contrasted with “instrumental value” or “instrumental worth”, the value that something has as a means to some other end. Intrinsic value does not need to rest on a thing’s intrinsic properties: something can be valuable-in-itself because of a relational property (e.g. rareness). “Intrinsic value”, then, is a bit of a misnomer, but I will continue to use it as it is less clumsy than *valuable-in-itself* or *final value* (a term that may lead to confusion for other reasons). Given the widespread technical use of “intrinsic value”, I trust my meaning is clear.

⁸ See Parfit (1984).

generations will enable greater well-being of future generations, and a greater total of well-being through history: the intuition that early sacrifice leads to greater pay-off. This effect would be amplified if there was an intuitive assumption that human history continues – that there is no future temporal limit such that the levels of well-being achieved at the end of the sequence above would end where they do.

Note that these factors are not relevant to the question as posed above. Crucially, the total amount of well-being throughout history – the *Historic Total of Well-Being* – is fixed. The question as posed stipulates a fixed number of people, evenly distributed throughout time, and assumes that the history of human well-being *ends*.⁹ Changing the temporal location of the well-being does not increase the Historic Total; there are no *instrumental* reasons for preferring any of the three distributions. In our examples, a sacrifice on the part of earlier generations can lead to later generations being better off, but not by increasing the Historic Total of Well-Being.

Our discussion in Part I will focus how well-being and its distribution contribute to the good, with a particular emphasis on Egalitarianism. Our goal is to evaluate how and whether Egalitarianism might be able to meet the challenges raised by application over time. The discussion will take the following shape: After explaining a few preliminaries, including my use of mathematical functions, focus on well-being, and proper temporal unit of comparison, I will proceed to present the Simple Temporal Case introduced above in greater detail, examining and rejecting potential sources of spurious intuitions. Next, I will lay out the Egalitarian View which we will be examining. We will then be in a position to evaluate how this principle responds to the temporal cases.

As will be discussed, the two most plausible egalitarian views to emerge in the

⁹ Indefinite extension of populations creates problems that I must leave for another discussion.

face of the Simple Temporal Case are Temporally Neutral Egalitarianism, which is indifferent between inequality at a time and over time, and Progressive Egalitarianism, which restricts an individual's complaint regarding inequality as being held relative to contemporaneous or past generations. We will then consider a more difficult series of cases, the Complex Temporal Cases, which place pressure on even application over time of the standardly-accepted Pigou-Dalton Condition.

Though the goal of Part I is not to flesh out or defend any particular theory, a new egalitarian position will emerge: Weighted Progressive Egalitarianism. The focus of the remainder of the thesis, in Parts II and III, is to evaluate whether this theory might be justified beyond its success in providing plausible judgments about the cases in Part I.

We will meet with some success in this endeavor, though we will find more questions than we do answers. However, as I shall argue, a plausible justification for the theory appears available, and the advantages that this theory stands poised to offer provide strong incentives to develop it further.

Part I

The Problem of Fairness & Time,

Part I

PRELIMINARIES

Before launching into our discussion, there are three preliminaries to address. The first concerns well-being; the second, the structure of the good and our use of mathematical functions; and the third, a simplifying assumption regarding the proper temporal unit of egalitarian concern.

2.1 PRELIMINARIES I: WELL-BEING

Distribution problems frequently begin as problems about resources. More specifically, they often begin with a problem about how to divide up some limited resource (in simple cases) or how to decide between a range of distributional possibilities involving a number of different, though still limited, resources (in more complicated cases). Underlying the problem of which distribution of resources is morally preferable are a number of factors, variously described as need, desert, and claims, which together form the basis for a process for distribution (determining who gets what) and the distributions themselves.¹ The problem is complex. At its heart is the rea-

¹ The relationship between moral factors like need, desert, and claims, will be explored further in Section II.

son we squabble over resources in the first place: resources are a means to making life go better, and a lack of resources can make life go worse. In short, resources contribute to well-being. In deciding distributions of resources, we are determining, in part, how well-off people are relative to each other, since we are deciding on distributions of a key determinant of well-being.²

One might then say: But distributions of resources also affect distributions of things like houses, and lifetimes; why not focus on these? The answer is that, again, the reason to care about houses and lifetimes is, in part, the goodness for an individual that can come from these things — that is, the well-being an individual might derive from them. Perhaps, there are other reasons to care about the distributions of such things,³ but it is clear that at least one of the reasons to care about these distributions is their potential to contribute to the well-being of individuals.

Our focus in this section, then, will be on distributions of well-being for two reasons: First, it might be thought to be a plausible starting point in determining a basis for evaluating distributions of resources or access to resources — or whatever the proper distributive object of moral concern is. This is because, as just described, well-being is a fundamental object of moral concern. That a resource might be used to make an individual's life go better is a reason to care about who gets it. Insofar as this is the case, before asking about which distributions of *resources* are better or worse, one must ask: does it matter how people fare relative to each other with respect to well-being?⁴ The answer to this question will help determine the answer

² Certainly, the mere distribution of resources is not the only factor which determines how individuals fare relative to one another. Other factors include how individuals choose to use those resources which they have access to, as well as factors unrelated to (material) resources, like one's attitude and relationships with other people.

³ For example, when individuals have certain types of claims, like property claims, to things like houses (if such claims are indeed possible), this might provide a reason to care about the distribution of houses. I will discuss claims further in Section II.

⁴ Would the better question be: does it matter how people fare relative to each other with respect to *opportunity* for well-being? Perhaps; but this seems — to me at least — to reflect the ideal that individuals have equal *claims* to well-being (discussed in Chapter 5), coupled with a recognition of the role that an individual, his attitude, and the choices he makes, plays in his own well-being. For this reason, I take equality of well-being as my starting point.

to the question about distribution of resources.

The second reason for beginning with a focus on distributions of well-being is that it simplifies some of the complexities associated with ranking distributions of resources. One of the ways in which it simplifies the problem is by reducing the type of distributive objects to one. When working with distributions of actual resources, this is not always possible, since which distribution of one particular resource is best might depend on how another resource is distributed. A further difficulty arises in ranking distributions of resources when the distributive or comparative principles must be applied over time, since there are a range of resource-types that appear as a result of time, including perishable resources (e.g. meat), resources easily shared by a number of people over time but not at a time (e.g. a house, property), recyclable and renewable resources, the creation of new resources (e.g. technology), and greater variation in the opportunities available to individuals and societies at different times (as contrasted with at different locations). I set these problems aside.

Well-being is a measure of how well-off (or not well-off) a person is. What makes up well-being — what constitutes well-being — is up for debate. So, too, is the question of what the *means* to well-being are; resources are just one example.⁵ I will not address these questions, but I hope that what I have to say will be applicable on most conceptions of the constitution of and means to well-being.⁶ The main point is that well-being is meant to serve as the measure for how well-off people are, thus making it possible to compare which lives go better or worse, and how individuals fare relative to one another.

But how should we measure well-being? I will assume a cardinal scale of well-being, but there are good reasons to be skeptical of the use of this measure. There

⁵ See Kagan (1998, 30-31).

⁶ I am partial to objective list theories. For a categorization of different positions, see Parfit (1984, Appendix I).

are many cases in which it might seem that the different possibilities for an individual's life cannot be compared. Unfortunately, defending this assumption against this concern is beyond the scope of this work, so I must take it for granted. But I can at least explain what it means.

It seems that there are many cases in which we can at least say that an individual's life goes better or worse: a life in constant pain is worse than one of pleasure. Insofar as we can say that one life would be better or worse for a person, we can use an ordinal scale. An *ordinal* scale is just a mathematical representation of a ranking: an ordering of the things in question. To illustrate, consider four places on a road out of Oxford: A, B, C, and D (Figure 2.1).⁷ We can then describe the location of




Figure 2.1: Road from Oxford

these places using an ordinal ranking, from first one on the road out of Oxford to the last, by assigning them each a number. The most obvious way to do so is to assign numbers 1 through 4 to the four places, but we might just as easily use any numbers that increase sequentially, like 2, 7, 9, and 28. Any four numbers will do, provided they preserve the order. This is because the numbers only represent the *order* of the places on the way out of Oxford; they offer close to no information about the relative distances between the places or between Oxford and the places. Similarly, an ordinal scale of well-being is just a ranking of lives for a given individual from better to worse. On an ordinal scale, to say that one life has well-being of 5 and another a well-being of 10 is simply to say the second life is better than the first (where the numbers increase with respect to increasing well-being).

⁷ I learned the Oxford distance example from Dan Halliday.

Consider, next, three possible ways a life would go: the first is a life of extreme pain owing to chronic illness, the second is a life of pleasure, health and happiness but with occasional headaches, and the third is the same as the second life but without the headaches. It seems plausible to say that the difference between the second life and the first life is much greater than the difference between the third life and the second life. The improvement from the first life to the second life is greater (removal of extreme pain) than the improvement from the second life to the third (removal of minor irritant). Insofar as we can compare the size of differences between our rankings, we have a *cardinal* scale. And it turns out to be a useful mathematical discovery that cardinal scales preserve the ratio of difference between points.⁸ So a cardinal scale measures *degrees* of well-being. One familiar example of a cardinal scale is in the use of degrees to measure temperature.⁹ The difference between 100 °C and 25 °C is three times the difference between 25 °C and 0 °C. Using our Oxford distance example, assigning $A = 4$, $B = 5$, and $C = 7$ is to say that the distance between B and C is twice the distance as between A and B. The numbers are just a scale that preserves the ratio of these differences; we could just as easily have assigned $A = 10$, $B = 20$, and $C = 40$, in the same way that we could just as easily have used degrees Fahrenheit to measure temperature instead of Celcius (the difference between 212 °F and 77 °F being three times the difference between 77 °F and 32 °F). As noted, I will assume this type of scale.

In order to get the discussion off the ground, however, we still need to be able to make these comparisons between people. To do so, we must be able to measure their well-being on the *same* scale. That is, we need *co-cardinality*. Co-cardinality is when the scaling factors of each individual's well-being are the same as the scaling factors of the others'. In other words, the utility function by which we rank and

⁸ Broome (1991).

⁹ Sen also uses this example (Sen, 1973, 3-4).

specify differences in well-being of particular lives must be the same for each person. This is, perhaps, a more dubious assumption, but defending it is beyond the scope of this work; I will take it for granted.¹⁰

2.2 PRELIMINARIES II: THE GOOD AND THE MATH

This thesis is about how well-being and different distributions of well-being contribute to the good — about how well-being affects the betterness of an outcome.

To this end, I make use of mathematical functions — “value functions” — to represent this ranking. To say that the total value of history (A), $V_{(A)}$, is 10, that the total value of history (B), $V_{(B)}$, is 20, and that the total value of history (C), $V_{(C)}$, is 40, is not to say that (A) has 10 “moral protons” and so forth. Rather, it is to say that (C) is better than (B) ($V_{(C)} > V_{(B)}$), and that (B) is better than (A) ($V_{(B)} > V_{(A)}$). Our question, then, is: how does the way in which well-being is distributed affect these rankings?

In addition to the betterness of an outcome, there is a betterness ranking of an outcome for an individual. This is that individual’s “personal good”.¹¹ Many theories take personal good to be synonymous with well-being. I, however, do not.

Throughout this discussion, an individual’s personal good is their well-being less any disvalue stemming from unfairness (as distinct from how that unfairness affects

¹⁰ Sen writes of this problem:

The question will be asked at this stage whether such comparisons are at all legitimate, and if so in what sense. Despite the widespread allergy to interpersonal comparisons among professional economists, it is I think fair to say that such comparisons can be given a precisely defined meaning. [...] If I say ‘I would prefer to be person *A* rather than person *B* in this situation’, I am indulging in an interpersonal comparison. While we do not really have the opportunity (or perhaps the misfortune, as the case may be) of in fact becoming *A* or *B*, we can think quite systematically about such a choice, and indeed we seem to make such comparisons frequently. (Sen, 1973, 14)

See also Broome (1991, esp. 215-218).

¹¹ It is also sometimes misleadingly called that individual’s “utility”.

how their life goes). Since we are discussing the value of equality and fairness, I want to leave open the possibility that a particular outcome might be better or worse for someone, even if it does not make them better or worse off in terms of well-being. Hence, when I say “make someone better off”, I mean in terms of their well-being: it makes their life go better. It is evident that it is better for an individual for their life to go better. The question is whether there is anything else in an outcome that might have value (or disvalue) for them. I will explain why this distinction shouldn’t be too counterintuitive in the context of introducing the value of equality.¹²

2.3 PRELIMINARIES III: TEMPORAL UNITS OF COMPARISON

There are important questions about what the proper temporal unit of concern is, lifetime or segments of time, in measuring well-being.¹³ These questions are particularly pressing in considering how comparative principles should operate over time: whether they should compare simultaneous segments, corresponding segments, whole lives, or some combination.

We set these difficulties aside. Unless otherwise stated, well-being values will be taken to be lifetime well-being, and our discussion will be restricted to distributions over non-overlapping generations. That is, individuals either have complete temporal overlap (are born and die at the same time) or none at all. As we shall see, there are enough problems with the simplified version — and attempting to resolve them may shed some light on how to approach the complications created by partial temporal overlap.¹⁴

¹² *Infra* 32.

¹³ *cf.* McKerlie (1989).

¹⁴ See Section 11.1, *infra* 230.

THE SIMPLE TEMPORAL CASE

We return now to the Simple Case and the Simple Temporal Case.¹

Earlier, I discussed the *prima facie* intuition in the Simple Case that (b) is better than either (a) or (c), and that there is no difference between (a) and (c) — they are the same distribution, differing only in *who* is better-off or worse-off, and not in the number of people at each level of well-being. There are a number of principles that could account for and underlie this position. The real trouble began, I suggested, not in considering more complex distributions, but in considering these same distributions over time — the Simple *Temporal* Case.

I noted that our intuitions changed when these were understood as distributions *over time* of well-being: (b*) no longer appeared the obvious best state of affairs, (a*) seemed less objectionable than (a), and — most strikingly — the indifference between the static versions (a) and (c) does not exist between (a*) and (c*), with (a*) appearing morally preferable to (c*). These *prima facie* intuitions about the Simple Temporal Case conflict with the Simple Case. But the only difference between the cases is that in the Simple Temporal Case different groups live at different times.

¹ *Supra* 1 and 3.

Does this justify such a radically different response to the Simple Temporal Case than to the Simple Case? Or are there other factors confusing our intuitions?

Begin with the comparison between (a^*) and (c^*) . *Prima facie*, (a^*) is better than (c^*) . Why? There seem to be two factors underlying this intuition, both linked to natural (though not uncontested) assumptions about time: growing populations and the question of humanity's survival, and an asymmetric causal relation. Two other factors — a deep epistemic problem posed by time and a temporal bias — may or may not affect our pre-theoretic judgments. Some of these we have already introduced. I will now consider each in greater detail.

3.1 SPURIOUS SOURCE I: GROWING POPULATIONS

First, intuitions about the Simple Temporal Case might be confused by assumptions about population size and the impact this has on the Historic Total of Well-being.

The first way this could occur is if we assume — contrary to the case stipulations — that the population extends beyond what is shown in the diagrams. This might result from a difficulty imagining human extinction *simpliciter*, or difficulty imagining that it ends where the case stipulates it does (the last generation of the shown distribution, whose lives end at time T). This might not be difficult in (c^*) , the series of generations with ever-decreasing well-being. But it might be particularly difficult to imagine human extinction to occur at the last generation of (a^*) , where the final generation is the best-off and marks the end of a series of generations with increasing well-being.

Should such difficulties occur, a preference for (a^*) would be natural if one assumed that the history of humanity continued on its upward trajectory: a indefinitely continuing series of generations with increasing well-being over time would lead to a greater historic total of well-being. Such a preference would be natural even

in cases where well-being reached a maximum and remained at that level indefinitely (Figure 3.1).




Figure 3.1: Indefinite End of Humanity

If one had difficulty imagining human extinction in the case of (c^*) as well, this would also explain a strong rejection of (c^*) : each generation becoming successively worse-off seems to imply that the end of humanity continues at a very low level of well-being, so that there are very many more people continuing to exist at such low levels. This intuition might become stronger depending on how worse-off such lives are — what exactly the “lowest” level of well-being means in real terms.

Note that neither of these two spurious intuition sources depends on difficulty imagining a *finite* number of generations. Rather, they turn on the assumption that human history continues beyond what is presented in the case at hand — that the distributions represent a trend that will continue, rather than the distribution over the whole of human history.

In either case, both the intuitive preference for (a^*) and rejection of (c^*) might be exacerbated by a further intuition about population and time: later generations are larger (have more people) than earlier generations.² This would have two effects: First, the historic total of well-being would be affected — larger in (a^*) and smaller in (c^*) . Second, this greater total would be achieved by having a greater number of better-off individuals in (a^*) and a greater number of worse-off individuals in (c^*) .³

² Of course, in the case of extinction, it might be reasonable to suppose that the final generations would be the smallest, as the human population tapers out. This would not explain a judgment against the distribution (c^*) , unless underlying this intuition is that, despite being few in numbers, their lives would be ones of extreme suffering and misery, not just a low level of well-being.

³ This is as contrasted with making the Historic Total larger by increasing or decreasing the

These would then not be symmetric distributions analogous to (a) and (c).

But the Simple Temporal Case stipulates against these reasons for favoring (a*) over (c*). First, the historic total of well-being is fixed: there is no difference in the *total amount* of well-being between (a*) and (c*). Second, each generation contains the same number of people — population-at-a-time is fixed. (a*) and (c*) are symmetric: there is the same distribution of well-being, with the same number of people living at each level. The difficulty of imagining this might confuse intuitions, but this is not, on reflection, a reason to prefer (a*) to (c*) in the Simple Temporal Case.⁴ Nor is it a reason to prefer (a*) to (b*).⁵

3.2 SPURIOUS SOURCE II: CAUSAL ASYMMETRY

The second factor that might affect our intuitions is a natural inference about a causal asymmetry between past and future generations. The direction of causation appears tied to the direction of time, at least intuitively: causes must be temporally prior to their effects. Though, in some special cases, causes might appear to be simultaneous, it is generally agreed that — as regards causation between events — temporally later events cannot cause temporally prior events to occur.⁶ If this is indeed the case, then there *is* an asymmetric causal relation between non-overlapping generations: past generations can affect future generations, but future generations

well-being of a single individual.

⁴ If need be, suppose human existence ends when the earth is bull-dozed for an intergalactic highway.

⁵ On an infinite — or unending sequence — it seems plausible that we would want to reach a sustainable level of well-being rather than oscillating indefinitely (e.g. $w_i(t) = \sin(t)$). Unfortunately, I must set aside the added complexity of an infinite or indefinite historic population.

⁶ For one discussion of this phenomenon, see Lewis (1986a); also Lewis (1986b). In some outlandish cases, like time travel, it is unclear how to appropriately describe the causation present, if any (Lewis, 1986c). Dummett argues that, while backwards causation seems implausible, its denial is more difficult to argue for than might appear (Dummett, 1964).

cannot affect past generations.⁷ Built into the structure of the Simple Temporal Case, then, are natural inferences about the causal relations possible between the better-off and the worse-off.

The difference in reaction to (a*) and (c*) might thus be linked to a natural inference about how such distributions came about — or what can be done about them. In (c*), for example, the higher level of well-being at earlier generations makes it appear that the lower levels of well-being at later generations — and perhaps even extinction itself — is a causal result of excessive levels of well-being in earlier generations. This is similar to the concern highlighted in recent days about excessive use of natural resources, low savings rates, and pollution. By contrast, in (a*), there is nothing that can be done by the better-off to fix the distribution. The worse-off generations have already past.

Could such asymmetries exist in a static case? Certainly — those who are better off seem far more well-equipped to affect the worse-off, while the worse-off, owing to their limited resources, are unable to affect the better-off. But these asymmetries seem to depend on *contingent* features of the state of affairs (e.g. societal structures or access to resources and transportation) and not on what appears to be some deeper metaphysical fact. The asymmetry in the temporal case, however, does not appear to be a contingent feature of society, but built into the ontological structure of our world. We might be able to remove the contingent causal asymmetry by changing our social institutions and access to resources. But we are unable to remove the causal asymmetry resulting from time. Given that this is a *deep* fact, it seems that our ethical thinking cannot ignore it.⁸

⁷ Indeed, this has often been used as a reason why the “conditions of justice” required by contractarian positions do not hold intergenerationally. There are good reasons to reject this argument. See Barry and Sikora (1978, 204-248)

⁸ What if this is not a deep fact about our world? Skepticism about causation, or time, might make the Simple Temporal Case collapse into a Simple Case: there, arguably, would be little difference between distributions over people at-a-time and distributions over people at different times. I must leave the normative implications of different metaphysical views on time for another discussion.

Either way, it has been stipulated in both the Simple Case and in the Simple Temporal Case that the inequalities have not been *caused* by any of the individuals. If we were able to accept this in the Simple Case, then we should also be able to conceive of it in the Simple Temporal Case. It may be a source of spurious intuition, but we should be able to set it aside.

This problem does, however, raise an interesting question: how does the presence or absence of causal relations between those within a distribution affect the goodness or badness of a distribution?⁹ A lack of concern for the inequality in (a*), given a deep fact about causal asymmetry, might be due to the inability of those better-off generations to do anything to make the earlier worse-off generations better off. Decreasing inequality in this history would thus require making later generations *worse off* rather than earlier generations better off (from the perspective of later generations). This seems to present an instance of a Temporal Version of the Leveling Down Objection: If there is nothing the better-off (later generations) can do to improve the lot of the worse-off (earlier generations), is there any respect in which it would be better for later generations to be just as badly off as earlier generations? (Figure 3.2) Likewise, in (c*), if the better-off *could* help the worse-off, and even

⁹ A similar concern appears in a slightly different context of the relevant population concept for holistic values. Arrhenius and Bykvist present the *Astronomy Conclusion*:

Whether it would be good or bad that a person comes into existence, will depend on facts about the welfare and the number of the people living billions of light years from here, even though this decision will not affect these people's welfare. (Arrhenius and Bykvist, 1995, 112)

They argue that the Astronomy Conclusion shows that the demarcation of a population in time is not enough. Rather, they suggest that the appropriate demarcation is captured by the idea of “community”, tentatively understood as “a set of people with continuous overlapping chains of interactions” (113-114). It is unclear that this is meant to track causal connections: they offer as one such appropriate demarcation under this view a demarcation between populations inhabiting the same planet who “have no contact whatsoever” and “who do not affect each other in any way”, like hermits and other isolated, indigenous people (113). Yet, there remains the possibility for strong causal connections between these groups: that they do not and have not affected each other does not mean that there is not the possibility of so doing. If such demarcations are meant to track the strength of causal connections, then the (strong) possibility of causal interaction would suggest a different conclusion than the one suggested by chains of interaction in this case.




Figure 3.2: Temporal Leveling Down

distributions are preferable, does this make what is already an undesirable situation worse?

In addition to this first question, there is a second interesting question raised by the problem of a “deep” causal asymmetry: what obligations exist between people who are either side of an asymmetric causal relation? And what strength do they have? That is, even if an axiology of distributions could be determined, how does it affect obligations to bring such distributions about for those who will never be able to reciprocate but who *are* affected by such decisions? This last question might be particularly difficult to address (or even raise) in the temporal case due to the Non-Identity Problem: since the decision about which distribution to bring about will affect who exists, it cannot be said of a given individual that they were *caused* to be made worse off by the choice of a particular distribution.¹⁰

Unfortunately, I must set these questions aside for now. I will return to the first one later.¹¹ The second is beyond the scope of this thesis.

¹⁰ For the original presentation and discussion of the Non-Identity Problem, see Parfit (1984, 351ff.).

¹¹ *Infra* 9.2.3.

3.3 SPURIOUS SOURCE III: EPISTEMIC CONSIDERATIONS

There remains a third factor that might be thought to affect our intuitions: a deep epistemic problem associated with time. We have good reason to be less certain in our beliefs about the further future. Small errors, when accumulated over time, result in large errors. “Punctuation events”, like the invention of the nuclear bomb or 9/11, have the potential to change the landscape in entirely unforeseen ways. As a result, there is a heightened level of uncertainty regarding events the farther into the future they might occur. This is a recognized difficulty in developing national strategy and rational life plans alike.¹²

But it would not seem that this is a source of spurious judgments in the case at hand. First, the case has been described from an atemporal perspective, using “earlier” and “later” generations, rather than “past” and “future”. The epistemic problem only arises from a particular perspective, namely, a future-looking one. Second, even if this might be readily forgotten — suppose that one had misread the case to be discussing these histories as starting *now* — it would be strange if this epistemic problem led to a preference of (a*) over (c*). Rather, if the epistemic problem did affect our judgments, it would be in the opposite manner: we have less information about the future, and so, it would seem, events in the distant future

¹² “Punctuation event” is a term used in strategic discussions to describe the kind of event that causes drastic changes to the status quo (e.g. 9/11). I first learned the term during a 2005 panel discussion with Jonathan Schell and John Gaddis on post-9/11 security strategy and nonproliferation.

This difficulty in planning is also recognized in philosophy. Rawls, for example, argues that uncertainty about the farther future is reason to have a hierarchy of plans and subplans:

A plan will, to be sure, make some provision for even the most distant future and for our death, but it becomes relatively less specific for later periods. Certain broad contingencies are insured against [...] but the details are filled in gradually as more information becomes available [...]. Indeed, one principle of rational choice is that of postponement [...]. We must not imagine that a rational plan is a detailed blueprint for action stretching over the whole course of life. It consists of a hierarchy of plans, the more specific subplans being filled in at an appropriate time. (Rawls, 1999, 360)

are less likely than events in the close future. Since it may be wise to discount for probability, this would lead to a greater valuing of temporally near events than temporally far, and (c^*) would be preferred over (a^*) . But I was not worried about a spurious source for *that* intuition, and so we may set this problem aside.¹³

3.4 SPURIOUS SOURCE IV: GROWTH OF RESOURCES & TECHNOLOGICAL PROGRESS

There may be a fourth possibility: it is a feature of our world that economic growth takes time. It is also a feature of growth that it requires sacrifice at earlier times in the form of abstention (from consuming a growing resource so that it may grow), labor, or technological development. Such growth provides an opportunity to affect the historic total of well-being: sacrifices made earlier can lead to a greater total overall.¹⁴ As a result, the badness of inequality resulting from growth might be thought to be mitigated by the goodness of the increase in well-being.

But this is not the case in the Simple Temporal Case as presented, for each history has the same Historic Total of Well-being. Still, one might think that (a^*) is not worse — not because the inequality is mitigated by the value of growth, but because the inequality does not matter in this case. Whether this view might be justified is one of the questions which this thesis attempts to answer.

¹³ Parfit provides a clear discussion of why discounting the value of future events for this reason — namely, that they seem to be less likely — is misguided (Parfit, 1984, 481-482).

¹⁴ Whether such growth is sustainable is a separate question.

3.5 SPURIOUS SOURCE V: PSYCHOLOGICAL PREFERENCE FOR PROGRESS

The fifth and final potential source of bias is a psychological preference for progress. It may be as a result of the correlation between progress (growth) and an increase in total well-being, or it may stem from some other feature of the human psyche. I point to it not as a certainty, but as a possibility. Intuitively, we find stories of overcoming adversity to rise to the top more pleasant than stories where someone at the top fell to the bottom. (Of course, such intuitions are affected by perceived desert of the hero or heroine. Villains meeting their demise seems somehow appropriate, or at least less bad.) Either way, there does seem to be a natural attraction towards things becoming better with time — not just that it is good that situations (lives, well-being, character) become better because they are then better, but because there is something good to the direction of the sequence from bad to good *over and above* the change in goodness of the situation.¹⁵ Likewise, there seems to be something bad about loss of the good — not just because it is good lost, but that the sequence from good to bad is worse *over and above* the change in goodness of the situation. The progression *matters to us* in a way that the distribution of the good over time has value in addition to the size of the total.

Whether we have such a bias — and, if so, why — is not my concern. As with a preference for growth, my concern is this: would a preference for progress be justified?¹⁶

¹⁵ *cf.* Kahneman et al. (1993).

¹⁶ If we do have such a bias, would this provide a debunking explanation for the intuition? The answer depends, to an extent, on the success of debunking arguments from psychological bias more generally. Provided that the existence of psychological bias does not undermine moral justification generally (a problem that is beyond the scope of the present work), the mere existence of a particular psychological bias does not debunk principles or intuitions consistent with it. Rather, we can ask whether those intuitive judgments are justified on independent moral grounds. This is the goal of the remainder of the thesis. Part I explores whether and how this intuition generates a problem for the egalitarian. Part II and Part III explores whether there might be independent justification for a theory consistent with the intuition. As will be discussed, it turns out that the

To begin answering this question, let us now turn to the application of distributive principles over time.

initial intuition that (a^*) is better than (b^*) seems to be picking up on the result of a slightly different principle: that comparisons between temporally distant individuals matter less than between those who are temporally close. Similarly, the intuition that (a^*) is better than (c^*) can be supported by principles unrelated to a value for progress.

CHAPTER 4

EQUALITY

Having set aside the spurious sources of intuitions about the Simple Temporal Case, we now turn to considering whether and how an egalitarian view might accommodate our intuitions. In the first two sections, I identify the type of egalitarianism on which our discussion will focus and explain how it determines the value of a distribution. In the third section, I consider how we might modify the axiology to accommodate the *prima facie* intuitions in the Simple Temporal Case. In the fourth section, I consider a series of more complicated temporal cases which suggest that a temporally neutral egalitarianism fails to capture what really matters about equality. Fortunately, a new view emerges which appears to address these concerns.

4.1 TYPES OF EGALITARIANISM

Egalitarianism is the name given to a range of views that hold that equality — of well-being, resources, or opportunity — matters. The underlying intuition uniting a variety of egalitarian positions is this:

it is bad, unfair or unjust, for some to be worse off than others through no fault of their own.¹

Call this the *Fundamental Egalitarian Insight*.

There are roughly two broad categories of egalitarianism, *deontological* and *teleological*. Deontological egalitarians object to inequality in a state of affairs when that inequality resulted from an unfair process. By contrast, they do not (usually) object to inequalities arising from fair processes. Parfit explains that, according to deontological egalitarians: “What is unjust, and therefore bad, is not strictly the state of affairs, but the way in which it was produced.”² This category of egalitarians also includes contractarian views of equality, like that of Rawls.³ I will set them aside. The focus of the current section is on the second type of egalitarianism, *teleological* egalitarianism, or *telic egalitarianism* for short. Telic egalitarianism holds that equality is itself *intrinsically* valuable, or valuable-in-itself.⁴ Undeserved inequality is bad, on this view, not because of how it came about but because inequality is bad in itself — that is, inequality detracts from the good of a state of affairs.

Telic egalitarianism thus places a value on there being equality within an outcome, and disvalue on inequality. We may therefore frame a discussion of egalitarianism in one of two ways: in terms of *equality* or in terms of *inequality*. Our discussion will proceed using the language of the latter, inequality.⁵ Hence, our question is about how to rank different inequalities.

Within egalitarianism, there is a distinction to be made between two types of egalitarianism: individualistic egalitarianism, and communal, or holistic, egalitari-

¹ Temkin (1993, 290).

² Parfit (1995, 90).

³ Parfit (1995, 90).

⁴ I discuss the use of intrinsic value, as contrasted with instrumental value, above on page 5.

⁵ The language of “inequality” is the language typically employed by economists. Even though the philosophical debate traditionally uses the language of “equality”, I follow the economists: the complexity of the problem is more easily captured in the language of reducing *inequalities*, rather than promoting *equalities*. See Temkin (1993, 7).

anism.⁶ Both positions favor states of affairs with less inequality over those with more inequality, but they differ in how inequality between individuals detracts from the value (or ranking) of a distribution.

Individualistic egalitarianism takes the good of equality — or more accurately, the bad of inequality — as a bad for the individual: those who are worse off have a *complaint* regarding inequality. A *complaint* is a measure of the unfairness suffered by an individual, whether or not they actually complain or would complain, based on a comparison between that individual and others in the state of affairs. It is a comparative measure, which takes the badness of that inequality for the person, above and beyond any instrumental badness resulting from the inequality (like jealousy, hurt feelings, or feelings of inadequacy).⁷

Communal egalitarianism, by contrast, takes the good of equality (or the bad of inequality) as a feature of the distribution or society. These positions are not just about equality, however. Indeed, the form of communal egalitarianism we will discuss — additively separable communal egalitarianism, otherwise known as prioritarianism — is not really a view about equality, since it is not concerned with *comparisons* between individuals. Rather, prioritarianism is the view that well-being has diminishing marginal value.

As noted earlier, we can express the way in which well-being contributes to the good — to the betterness orderings of outcomes — using value functions.⁸ Taking a moment to discuss the formalization of these positions will help clarify exactly how these positions differ.

⁶ Hereafter, ‘egalitarianism’ refers to telic egalitarianism.

⁷ The term is also not meant to imply that the better-off are not concerned with inequality or that they would not complain about their situation as being better off.

I follow Temkin in the use of this term. As he writes, the idea behind the term is this:

for any situation where some people are better off than others, we can say that the best-off have nothing to complain about while the worst-off have the most to complain about [regarding inequality]. (Temkin, 1993, 19)

⁸ Section 2.2, page 14.

The *contributive value* of an individual is the value that their personal good contributes to the overall value of a state of affairs of which they are a member. That is, an individual's contributive value is a function of their personal good:

$$v_i = v(g_i) \quad (4.1)$$

where v_i is the contributive value of person p_i , and g_i is his personal good.⁹

The total value of a distribution, (A), is a function, V , of the contributive values of the distribution's members. Since, on this model, the contributive value of an individual is a function which takes as its only inputs the personal good of that individual, V can be expressed as a function of the personal good of the individuals:

$$V_{(A)} = V(v_1, \dots, v_n) = V(g_1, \dots, g_n) \quad (4.2)$$

What form does V take? For purposes of our discussion, I am going to assume that V is additive for fixed (same number) populations.¹⁰ That is:

$$V = \sum_{i=1}^n v_i(g_i) \quad (4.3)$$

⁹ I use “v” as shorthand for “value”, and the subscript, i , to denote the individual person p_i . The term “contributive value” is from Arrhenius (2000, 7).

¹⁰ My reasons for doing so have to do with assumptions defended strongly in Broome (1991), namely, coherence of the individual and general betterness relations, and the Principle of Personal Good:

- Principle of Personal Good** (a) Two alternatives are equally good if they are equally good for each person.
- (b) If one alternative is at least as good as another for everyone and definitely better for someone, then it is better.

This is the claim that the value of a fixed population depends solely on the value of individual good and is increasing in this value.

Note that communal egalitarianism, except for additively separable communal egalitarianism (“Prioritarianism”), violates the Principle of Personal Good. I will not argue against those views here, but refer the reader to Broome (1991, Ch.9).

I can now make clear the difference between individualistic egalitarianism and additively separable communal egalitarianism.

Individualistic egalitarianism takes an individual's personal good to be the difference between the *noncomparative* value of their well-being and the disvalue of any individual *comparative* complaints. This may be given:

$$g_i = w_i - \alpha c_i(w) \quad (4.4)$$

where w_i is p_i 's well-being; c_i is p_i 's complaint regarding inequality (hereafter, 'complaint'), which is a function of p_i 's well-being w_i and the well-being of the other individuals in the state of affairs, denoted by the vector $w = \langle w_1, w_2, \dots, w_n \rangle$; ¹¹ and α is a scaling factor that affects how much weight is given to the disvalue of comparative complaints relative to the value given to well-being considered noncomparatively. To illustrate, consider a simple world with only two people, p_i and p_j , who have well-being w_i and w_j respectively. g_i , the personal good of p_i , may be expressed as:

$$g_i = w_i - \alpha c_i(w_i, w_j) \quad (4.5)$$

All this is a formal way of saying that an individual's personal good (i.e. the individual betterness relation) is determined not only by their own well-being, but also the well-being of those others in their world. Note that this entails that a situation can be worse for someone even though it does not make his life go worse. Depending on the form and weight given to complaints, this entails that making the best-off better off without making the worse-off worse off (in terms of well-being)

¹¹ Recall that a vector is just a list.

might not increase the goodness of the distribution because it is not the case that the good of the worse off has stayed the same (they have an increased complaint). This shouldn't be too counterintuitive: giving an extra pool or Prada bag to a real housewife of Orange County makes her better off without making a starving kid's life go worse. But it seems pretty obvious that it doesn't make the situation any better, and also that it increases the unfairness that the kid suffers (even though he's not aware of it).

The exact way in which the complaint is determined — i.e. the form of c_i — is a matter for debate. In more complicated cases, the domain of c_i — whose well-being is included on the list against which p_i 's well-being is compared — is also contentious. When p_i has no complaint (e.g. p_i is not worse off than p_j), then $c_i(w_i, w_j) = 0$. An example of p_i 's complaint in a simple, two-person world is:¹²

$$c_i = \max(0, w_j - w_i) \quad (4.6)$$

Equation (4.6) expresses the following intuitive egalitarian idea: If p_i is equally as well-off as or better off than p_j then he has no complaint ($c_i = 0$). If p_i is worse off than p_j , then his complaint is the size of the gap between his well-being and that of p_j .

By contrast, communal egalitarianism takes the bad of inequality as a feature of society, state of affairs, or fixed distribution rather than as a feature of an individual's good. Communal egalitarianism has two types, one additively separable (the total value of a distribution is the sum of individual contributive values, as in (4.3)) and one that is not (the total value of the distribution takes some other form). I have already set aside the latter by assuming that the general value function is additive.

¹²I borrow this example from Broome (1991, 182).

Additively separable egalitarianism, as noted, is also called Prioritarianism. I will use “Prioritarianism” for it hereafter, since “egalitarianism” is something of a misnomer.¹³

There are two types of Prioritarian. One weights the contributive value of an individual’s personal good according to how much personal good that individual has: the more good an individual has, the less each additional unit of good counts towards the general good:

$$v_i = v(g_i) \tag{4.7}$$

where $v(\cdot)$ is a strictly concave transform. This view might be described as “diminishing marginal contributive value of personal good”.

The second type focuses directly on an individual’s well-being. The second type might be described as “diminishing marginal good of well-being”: the more well-being an individual has, the less each additional unit of well-being counts towards personal good.

$$g_i = g(w_i) \tag{4.8}$$

where $g(\cdot)$ is a strictly concave transform.

As can be seen, prioritarianism is essentially *noncomparative*. It is a view about how an *individual*’s well-being contributes to either their good, or to the betterness of an outcome. There is nothing in here that involves a comparison. Hence, even though it may be said to aim at equality — since the total value function will be maximized when good, or well-being, is distributed equally — it is not properly considered an egalitarian position.¹⁴

¹³ The term “Prioritarianism” is taken from the name which Parfit uses for a similar and related view, the “Priority View” (Parfit, 1995). Broome calls the view “additively separable communal egalitarianism” (Broome, 1991). Temkin calls this view “extended humanitarianism” (Temkin, 1993, 245-246).

¹⁴ This argument has been made by Broome, Parfit, and Temkin. (Broome, 1991; Parfit, 1995; and Temkin, 1993, Ch. 9.)

Since Prioritarianism is noncomparative, I set it aside. Instead, our discussion will focus on individualistic egalitarianism, hereafter simply “egalitarianism”.¹⁵

Before continuing, there is one final point to make. Since an individual’s personal good on the views discussed takes as its only inputs the well-being of that individual and the well-beings of the others in the distribution, V can also be expressed as a function of the well-being of the individuals:

$$V_{(A)} = V(w_1, \dots, w_n) \quad (4.9)$$

4.2 COMPLAINTS, AND MEASURING THE DISVALUE OF INEQUALITY

Individualistic egalitarianism, however, still encompasses a range of views, many plausible, that have deep differences in measuring the badness of inequalities despite apparent overlap in the more frequently discussed — and simpler — cases.¹⁶ Within individualistic egalitarianism, there is a further question about what form the complaint, c_i , should take. In particular, three questions must be answered:¹⁷

1. *Who* has a complaint?
2. What are their complaints relative *to*?

¹⁵ What if equality is not what matters? What if what we really care about — as I tend to think — is that people get what they deserve, and that equally deserving people get equally? (For a defense of this view, see Kagan (1999).)

How this position will be formalized depends on whether absolute desert matters, or only comparative desert, and also on whether offenses against desert are a matter of personal or general good.

¹⁶ See Temkin (1993) for an excellent discussion, especially Ch. 2-3. The question of the appropriate way to measure an inequality arises in both a descriptive (which distribution is *more* unequal) and normative (which inequality is *worse*) context. These two contexts should not be confused; my concern is with the normative aspect of inequality. For a discussion of the difference between the two contexts, and various ways of measuring inequality (though not always from an individualistic egalitarian perspective) see Sen (1973).

¹⁷ I here follow Temkin’s lead.

3. *How* are these complaints to be weighted?

These questions do not have easy answers, even in the static case. I will briefly explain some of the more popular options. My goal is not to defend a particular static view, but to lay the groundwork for the discussion which comes next about the implications of applying these views over time.

The first question, *who* has a complaint regarding inequality, is about for whom c_i has some positive value, and for whom $c_i = 0$. It has two main contending answers: those who have less than the average, and all but the best-off.¹⁸ On one of these views, the less-than-average view, only those who have less well-being than the average have a complaint. All those who have more well-being than or the same well-being as the average have no complaint ($c_i = 0$). On the other view, the all-but-the-best-off view, everyone has a complaint except the best-off. That is, for all but the best-off, c_i has some positive value.

The second question, what are their complaints relative *to*, is about the *scope* of an individual's complaint. The answer to it determines what form the complaint, c_i , takes. There are several standard answers: complaints might be determined, in part, by comparing the well-being of the individual to that of the best-off person (BOP), to the average (AVE), or relative to all those better off (ATBO). For example, under BOP, a complaint will be a function of the difference between the best-off person and the person with the complaint:

$$c_i(w) = c(\max(w) - w_i) \quad (4.10)$$

¹⁸Temkin (1993, 25) Another possible answer is “the worst-off”. This appears to be the answer given by Maximin, but this is actually confusing the answer to first question (who has a complaint?) with the answer to the third question (how do we weight the complaints?). Those who are not the worst-off may have a complaint regarding inequality, but a principle of equality need not give it much (if any) weight.

and might take the following form:

$$c_i(w) = \omega_i * (\max(0, \max(w) - w_i)) \quad (4.11)$$

where ω_i is a scaling factor that affects the size of the complaint — that is, how much this comparison matters normatively. I call this comparison the *complaint base*. In (4.10) and (4.11), the complaint base is $(\max(w) - w_i)$.

Note that if p_i is the best-off person, then he will not have a complaint ($c_i = 0$ where $\max(w) = w_i$); if p_i is not the best-off, then he will have a complaint sized in proportion to how badly off he is as compared to the best-off person.¹⁹

The third question, *how* are the complaints to be weighted, is about how bad the relevant comparison is (whatever it was measured against). This is to ask what form $c(\cdot)$ takes.

In addition to these questions about the complaint, there is a question about how much this disvalue matters relative to well-being — that is, what is α ? An additive principle (AP) weights all complaints the same ($\alpha = \text{constant}$). A weighted additive principle (WAP) weights individuals' complaints according to how well-off they are (either noncomparatively or relatively speaking). Under this view, the weighting factor, α , is a function of either the individual's well-being (where the weight is determined noncomparatively) or of the individual's well-being and that of the others in the state of affairs (where the weight is determined comparatively). Assuming nonnegative well-being, an example of a WAP scaling factor that weights the complaints of the worse-off more than those of the best-off is:

$$\alpha = \frac{1}{w_i} \quad (4.12)$$

¹⁹ For a discussion of the various merits and demerits of these views, see Temkin, 1993, Ch.4, also p. 285. Temkin favors BOP on the grounds that its “spirit and underlying reasoning are important elements of inequality that must be accounted for” (Temkin, 1993, 285).

The maximin principle (MP) is an extreme version WAP, which only gives weight to the complaint base of the worst-off.²⁰

Armed with a clear (or, at least, *clearer*) understanding of the various ways of measuring static inequality, we turn to the problems created by time.

4.3 EQUALITY, THE SIMPLE CASE, & THE SIMPLE TEMPORAL CASE

Return to the Simple Case. We can now see more clearly how egalitarianism favors (b) over (a) and (c), but is indifferent between (a) and (c).

Remember that in the Simple Case, there are k groups of n people each, with each person within a group having the same well-being as everyone else in their group.²¹ Notate individuals as $p_{i,j}$ where i specifies the group ($i = 1, 2, 3, \dots, k$) and j specifies the individual within the group ($j = 1, 2, 3, \dots, n$). The vector (list) representing the i th group is given as G_i , where $G_i = \langle p_{i,1}, p_{i,2}, p_{i,3}, \dots, p_{i,n} \rangle$. All the functions explained above remain the same, but with subscript “ i, j ” instead of “ i ”.²² The contributive value of an individual is still given as his personal good:

$$v_{i,j} = g_{i,j} = w_{i,j} - \alpha c_{i,j}(w) \quad (4.13)$$

where w is the vector of the well-beings ($w = \langle w_{1,1}, w_{1,2}, \dots, w_{i,j-1}, w_{i,j+1}, \dots, w_{k,n} \rangle$).

The total value of a given state of affairs is given as a sum of the individual

²⁰ Temkin offers a few additional principles of equality, but I set them aside. *cf.* Temkin (1993, Ch. 2-3); Sen (1973).

²¹ I used groups, even in the simple case, so it is directly parallel to the Simple Temporal Case, where generations replace groups.

²² The subscript is just a means of identifying the individual.

values:

$$V = \sum_{i=1}^k \sum_{j=1}^n v_{i,j} \quad (4.14)$$

For a given total of well-being, the total value is maximized when there is no disvalue stemming from complaints of inequality. The distribution in (b) does this: Since everyone has equal well-being, no one has a complaint of inequality, and so $c_{i,j} = 0$ for all i and all j . As a result, in distribution (b), $v_{i,j}$ is maximized for all i and all j , and so V is maximized. Each of the principles of equality noted above agree in this judgment, that (b) is the best distribution.

The equality principles outlined above also offer the same evaluation of (a) and (c). This is because (a) and (c) are simply permutations of each other: G_1 in (a) is equivalent to G_k in (c), G_2 in (a) is equivalent to G_{k-1} in (c), and so forth. As a result, regardless of how the complaints are determined, they will likewise be permutations of each other: the complaint of an individual in the first group of (a), $c_{1,j}$, will be equal to the complaint of an individual in the last group of (c), $c_{k,j}$, the complaint of an individual in the second group of (a), $c_{2,j}$, will be equal to the complaint of an individual in the second to last group of (c), $c_{k-1,j}$, and so forth.²³ Recall that since V is determined by summation, the total effect of individual complaints on the goodness of a distribution of well-being is the sum of the individual complaints. That is, V may be re-written as the difference between the sum of the noncomparative components of $v_{i,j}$ and the sum of the weighted complaints regarding inequality:

$$V = \sum_{i=1}^k \sum_{j=1}^n w_{i,j} - \sum_{i=1}^k \sum_{j=1}^n \alpha c_{i,j} \quad (4.15)$$

The same complaints are present in both (a) and (c) — they just belong to different people — and there is the same total amount of well-being (hence the moral good-

²³ This is also true for weighted complaints since the value of the scaling factor, α , will similarly be permuted.

ness of well-being determined noncomparatively is the same). Thus, *any* egalitarian principle outlined above will be indifferent between (a) and (c), and will rank them both as worse than (b).²⁴

Now consider the Simple Temporal Case. As presented above,²⁵ the k contemporaneous groups G_i are replaced by k *non-overlapping* generations g_i .²⁶ How should an egalitarian respond to the Simple Temporal Case? That is, how should egalitarian principles apply over time?

One answer is to argue that time makes no difference: location in time is not a morally relevant factor in determining the goodness of distributions any more than location in space.²⁷ Call this *Temporally Neutral Egalitarianism (TNE)*. Temporally neutral egalitarianism, on its face, seems plausible: if we care about equality between persons, we care about equality between all persons. Recall that the fundamental insight underlying egalitarianism is that it is bad for people to be worse-off than others through no fault of their own.²⁸ An individual has no control over their temporal location.²⁹ This being the case, it seems at least plausible — if not required — that the fundamental insight be applied across times. In such a way, the temporally

²⁴ Coincidentally, communal egalitarian principles will also yield this judgment of indifference between (a) and (b).

²⁵ Section 1, page 3

²⁶ Individuals within the same group or generation have complete temporal overlap. The k groups may be said to have *complete temporal overlap*. By contrast, the k generations have *no* temporal overlap. Thus, individuals in the Simple Temporal Case either have complete temporal overlap or none at all. There is an interesting debate concerning situations in which there is partial temporal overlap which I set aside for now. We will, however, return to the importance of this in closing (Section 11.1, page 230). For a more detailed discussion of these issues, see McKerlie, 1989.

²⁷ Of course, there are some situations in which *mere* spatial distance *may* be thought to make a moral difference. For an interesting new discussion about how distance might be relevant to the duty to aid, see Kamm (2007, Ch. 11-12).

For a discussion of this question in relation to population demarcation and the suggestion that space — or more accurately, community membership — is relevant, see Arrhenius and Bykvist (1995, 67, 112ff.). For the view that demarcating for time forces demarcating for space to avoid arbitrariness, see Carlson (1998, 290-291).

²⁸ Page 28.

²⁹ This is true with two exceptions which I set aside: suicide, by which a person limits the times they extend over, and time travel.

neutral egalitarian response to the Simple Temporal Case would parallel the Simple Case: groups G_i are replaced by non-overlapping generations g_i without affecting any of the equations. $v_{i,j}$ takes the form given by (4.13) above, and V takes the form given by (4.14). TNE thus ranks (a^*) , (b^*) , and (c^*) in the same way as (a) , (b) , and (c) : preference for (b^*) over both (a^*) and (c^*) , and indifference between (a^*) and (c^*) .

This may be too quick. These judgments do not coincide with our pre-theoretic intuitions about the Simple Temporal Case. In particular, they seem to not leave any room for our pre-theoretic judgment that (a^*) is preferable to (c^*) . While we have yet to determine whether and under what conditions such a temporal bias is justified, a new question arises about whether an egalitarian view can even make room for the temporal bias. I will return to this in a moment.

First, flag for now another question that is perhaps a more obvious concern with TNE: is it really plausible for individuals to have complaints (regarding inequality) against other times? This is a question about the scope of complaints — what are complaints measured relative *to*. But it might be phrased as a question about the form complaints take — are complaints against the better-off in other generations to be given the same weight as complaints against the better-off in one's own generation? Evaluating this will require considering more complex cases, contrasting inequalities within a generation against inequalities across generations. I will address this in the next section.³⁰

Before moving on to the complex cases, however, we must address the question of whether a temporal bias — justified or not — is even compatible with egalitarianism. There are three ways we can accommodate or explain this bias, each reflected in alterations to the above equations:³¹ by weighting future well-being (altering $v_{i,j}^*$),

³⁰Section 4.4, page 53.

³¹ Perhaps, there are other ways to factor it in. I am unaware of them. Alternatively, perhaps

by accommodating a preference for progress (altering V), or by questioning the scope of individual complaints regarding inequality (altering $c_{i,j}$). I will discuss each in turn. As we shall see, not all are equally plausible.

4.3.1 ADJUSTING $v_{i,j}$: WEIGHTING FUTURE WELL-BEING

The first way to factor in a temporal bias is to alter $v_{i,j}$ so that the value of an individual's well-being, $w_{i,j}$, varies according to where they are located temporally. $v_{i,j}$ becomes a function of both $w_{i,j}$ and time, t . For simplicity, assume that the goodness of well-being varies according to which generation an individual belongs. (In this simplified version, $v_{i,j}$ could be written instead as a function of $w_{i,j}$ and i .) Call this type of temporal bias *Temporally Variable Noncomparative Value of Well-being*, or TVV*.

In our *prima facie* response to the Simple Temporal Case, the bias was a bias towards the future. Future-biased TVV* would then hold that:

F* An individual's well-being has greater value (considered noncomparatively) the later in time that individual lives (the later the generation of which that individual is a part).³²

In other words, TVV* discounts, or weights, the value of well-being by its location in time. In Future-biased TVV*, the value of an individual's well-being is discounted more severely the earlier the generation of which they are a part. This is the opposite of a type of temporal discounting of the value of future well-being, called "pure social discounting", which has generally been rejected by philosophers.³³ This does not bode well for future-biased TVV*.

the temporal bias would appear as a separate factor or serve a tie-breaking role. On either of these latter two views, this appears to me to be part of a question about the form of V .

³² N.B. This differs from saying that time is a morally relevant *difference* in considering comparative differences between people.

³³ *cf.* Broome (1992, Ch.3); Cowen and Parfit (1992); Parfit (1984, Appendix F).

Is our future temporal bias *really* that the noncomparative *value* of well-being increases over time? This seems strange. Consider a simplified version of Cowen's argument from displacement of persons against discounting over time.³⁴ Mary, an ordinary human being, lives in 2009. She has a well-being of 10 as has everyone who ever existed or will exist in the history of her world.³⁵ What if we were able to move Mary in time without affecting anything else (a far-fetched assumption, but this is a far-fetched case)? Suppose Mary lived in 2109 instead. The rest of the world remains exactly as is: Mary has well-being 10, as does everyone else, and the total historic population of the world remains the same. The only change, then, is that the 2009 population has been decreased by one, and the 2109 population has been increased by one. Would this be an improvement in any way?

It is hard to see why it would be. But that is precisely what future-biased TVV* insists: well-being matters more (morally speaking) the later in time the person who enjoys it lives. Future-biased TVV* prefers the world where Mary lives in 2109 to the world where Mary lives in 2009. Yet, the only difference between the two worlds is where in time Mary (and, consequently, her well-being) is located. That one scenario is better than the other does not seem plausible — and neither does TVV*.³⁶

Note that rejecting TVV* is *not* the same as denying that time might be morally relevant. Rather, it is denying that the value of an individual's well-being considered *non-comparatively* varies according to where in time they are located. As we shall see, time might be a morally relevant factor *comparatively*-speaking.

So, a future bias of the sort that places greater *value* on the well-being of future

³⁴ Cowen (1992) This argument was originally made against past-biased TVV*, but also applies to future-biased TVV*.

³⁵ The assumption that everyone in Mary's world has equal well-being is designed to help remove any confusions about value stemming from inequalities. What if not everyone had level well-being 10? We will discuss this later (*infra* 62).

³⁶ Not only does the judgment that one scenario is better than the other seem implausible, consequentialism implies that such a judgment is false. For the full argument, see Cowen, 1992.

generations is not justified. A different justification must be found if we are to maintain the judgment that (c^*) is more objectionable than (a^*) , or that the difference between the value of (b) and (a) is larger than the difference between the value of (b^*) and (a^*) .

4.3.2 ADJUSTING V : ACCOMMODATING A PREFERENCE FOR PROGRESS?

One explanation for the temporal bias might be a preference for progress. That is, what we really care about is progress: we think it makes things better, all else being equal, when things get better over time. This is different from the view just presented, future-biased TVV*, in that it does not discount the non-comparative value of well-being depending on location in time. Rather, what a preference for progress does is attach a value to *change* in well-being between times, and, in particular, increases in well-being. Call this general view *Progressive Welfarism*.³⁷

One way to formulate this view is to take progress as value added to the goodness of a particular distribution through time. This value is given by $P(w_{1,1}, \dots, w_{k,n})$, which is a function over the well-being of all those who ever exist. P is positive when progress adds to the total value of a particular historic distribution, and 0 or negative otherwise. The total value of a historic distribution of well-being is then given by:

$$V = \left(\sum_{i=1}^k \sum_{j=1}^n v_{i,j} \right) + P(w_{1,1}, w_{1,2}, \dots, w_{k,n}) \quad (4.16)$$

where P is the intrinsic value of progress over time.

There seems to be a conflict, however, between egalitarian sensibilities and valuing progress *per se*. Progressive welfarism values worlds with *greater* inequalities between generations. Progressive welfarism favors sacrifice and deprivation of some,

³⁷ Whether welfarism can include other values is a point of debate. By “progressive welfarism”, I mean the view that whether well-being increases over time matters in addition to the added value from the well-being itself — that the *sequence* of well-being levels contributes or detracts from the value of a historic distribution of well-being.

not for the sake of increasing overall well-being, but for the sake of making those who exist later better off compared to those who live at earlier times. Valuing progress in this way doesn't just mean that there is a sense in which it is *not worse* for some to be worse off than others through no fault of their own. Valuing progress in this way means that there is a sense in which it is *better* for some to be worse off than others through no fault of their own — namely, when they come earlier in the temporal sequence. This runs directly counter to the Fundamental Egalitarian Insight, which states that it's a bad thing for some to be worse off than others through no fault of their own.³⁸

Is this just a case of ordinary conflicting values? After all, one state of affairs may be better than another with respect to one aspect of value, like well-being, but worse with respect to a different aspect, like equality. But valuing progress while maintaining a temporally neutral egalitarianism is at odds in a way that valuing increasing well-being and equality of well-being is not: progressive welfarism and temporally neutral egalitarianism, in evaluating a state of affairs where later individuals are better off than earlier individuals, are evaluating the same aspect of the state of affairs — the *comparison* of well-being between those individuals. One says that the inequality in this comparison is, in itself, good; the other says it is not. They conflict on the judgment of the same aspect. This is quite different from the conflict between valuing well-being and egalitarianism, where the difference in judgment stems from evaluating different aspects.

A different variety of egalitarianism which is not temporally neutral might be able to accommodate this value. This would be the case if such a version of egalitarianism could avoid saying that the inequality that favors later individuals over earlier individuals (through no fault of their own) is bad. Such a view would not come into conflict with progressive welfarism, which is the view that such an inequality is good.

³⁸Above, page 28.

And so, the hunt for a temporal bias consistent with egalitarianism continues.

4.3.3 ADJUSTING $c_{i,j}$: QUESTIONING THE SCOPE OF INDIVIDUAL COMPLAINTS

Consider the three distributions in Figure 4.1.³⁹:




Figure 4.1: Temporal Leveling Down Objection Revisited

In each distribution, well-being increases steadily from g_1 to g_{present} . In (d*), the increase continues at the same rate from g_{present} to g_k . In (e*), well-being stays at the same level from g_{present} to g_k . The historic total of well-being is thus less in (e*) than it is in (d*). This might lead us to wonder if there is *any* respect in which (e*) is better than (d*). That is, is there any respect in which it would be better for future generations to continue at our present level of well-being instead of enjoying higher levels of well-being unavailable to us? The *Temporal Leveling Down Objection* answers No. For comparison, the traditional (Static) Leveling Down Objection is pictured in Figure 4.2.

³⁹ These distributions are similar to those presented earlier in Figure 3.2 (page 22).




Figure 4.2: Static Leveling Down Objection

According to TNE, there is at least *something* to be said for (e*) over (d*) even though the historic total of well-being is drastically reduced: namely, there will be fewer complaints regarding inequality.⁴⁰ This is *not* to say that TNE is committed to saying that (e*) is better than (d*) all things considered, but to the position that there is at least *one* respect in which (e*) is better than (d*), namely inequality.

It seems difficult to accept that this respect matters in this case.⁴¹ The Temporal Leveling Down Objection packs with it greater force than the (Static) Leveling Down Objection: how could it be better in any way for future people to have well-being at our level rather than more? There may be greater equality in (e*), but this equality has no value. Suppose we concede: the *pro tanto* reason in favor of (e*) stemming from equality is small, but existent. Even so, it seems that the difference in moral value between (d*) and (e*) is at least somewhat greater than the difference in moral value between (d) and (e). More strongly still, it is hard to see that the

⁴⁰ This will be true with the exception of Maximin (MP). According to a temporally neutral MP egalitarian, (d) and (e) are equally bad with respect to inequality.

⁴¹ Those who find the (Static) Leveling Down Objection convincing, and hold that there is *no* respect in which the static equivalent of (e) is better than (d) — that is, those who are not egalitarians — will likely have anticipated the Temporal Leveling Down Objection. They are not our present concern, however. Our question is how egalitarians — those who *do* value equality in the static cases — should apply their principles over time.

greater equality in (e^*) provides a reason of similar, if any, force in favor of (e^*) over (d^*) as compared with the reason that the greater equality in (e) provides to favor (e) over (d) . In other words, it is hard to see the offense of inequality in (d) and (d^*) as equally bad (or, alternatively, the greater equality in (e) and (e^*) as equally good).

This leads to two questions: Can we consistently hold that, while the greater equality in (e^*) and (e) provide (*pro tanto*) reasons in their favor, the reasons do not have the same strength? More strongly, can we consistently hold that there is a way in which (e) is better than (d) without holding that there is *any* way in which (e^*) is better than (d^*) ?

There might be a way to do so. An option remains, namely, altering the scope and form of the complaint. While it seems plausible that individuals could have a complaint regarding inequality relative to those who live at the same time, it seems less plausible that individuals could have a complaint relative to those who are temporally distant. Is there really anything *bad* about inequality between very temporally distant generations? If so, it seems that such badness would be negligible — or at least less bad than similar inequalities between temporally closer generations or contemporaneous people. I will explore this question later in this chapter.

For now, allow me to make an observation about one version of temporal bias that appears to be compatible with egalitarianism. This temporal bias constrains who complaints regarding inequality might be *to* by restricting the domain of $c_{i,j}$. For simplicity, suppose that there are no alterations to the *form* of $c_{i,j}$ in the sense that temporal distance does not factor into the weighting of individual complaints. Rather, the domain of $c_{i,j}$ is restricted to the well-being of those past and present. That is, individuals cannot have complaints regarding inequality relative to the members of later generations, but have equally weighted complaints of inequality relative to the members of prior generations and the contemporaneous generation.

One way of achieving this is for the complaint to take the form:

$$c_{i,j} = c(\max(w_{1,1}, \dots, w_{i,n}) - w_{i,j}) \quad (4.17)$$

This is a very simple model, but will serve our present purposes. The complaint given by equation (4.17) measures complaints relative to the best-off person (BOP) located at the same time as or at an earlier time than the individual whose complaint it is. Call the restriction of complaints as being relative to prior and contemporaneous generations the *Progressive Scope Restriction*. Call this family of egalitarianism that restricts in this manner *Progressive Egalitarianism (PE)*.

Return to our cases. Consider first the Simple Temporal Case. Since there are only complaints of inequalities relative to contemporaneous or past generations, there are only complaints of inequality in (c*). That is to say, there are *no* complaints of inequality not just in (b*), but also in (a*). Given that the historic total of well-being is fixed between the three histories, this means that the total value of (a*) is the same as of (b*), both of which are greater than the value of (c*) ($V(a^*) = V(b^*) > V(c^*)$). PE is thus compatible with the type of temporal bias found in our *prima facie* reaction to the Simple Temporal Case.

Consider now the Temporal Leveling Down Objection Revisted. None of (d*), (e*), or (f*) contain complaints of inequality. The total historic amount of well-being is fixed between (d*) and (f*), and is less in (e*). PE is indifferent between (d*) and (f*) ($V(d^*) = V(f^*)$), but is able to rebut the Temporal Leveling Down Objection: in (e*) there are *no* complaints regarding inequality. As a result, there is *nothing* to be said in favor of (e*) over (d*). This is the stronger version of the bias described above.

PE also captures a basic intuition regarding three additional questions: Would it be better if past generations had been better off than they were? Yes, by increasing the amount of well-being. Would it still be better if this put us, g_{present} , at a lower

level of well-being than we currently are? Maybe, depending on how well-being is distributed, and if the total well-being is not thereby decreased. Would it still be better if this put us at a lower level of well-being than they were? No, not unless the gains to their well-being exceeded the disvalue of our complaint regarding inequality.

These are interesting results. But we might worry: is PE *too* simplistic? Consider the following four distributions: Sine-step, Cosine-step, Level, and Up-step. All contain the same population — 10 generations of n people each — and the same historic total of well-being. How do they compare on complaints of inequality?




Figure 4.3: Sine-step

In Sine-step (Figure 4.3), the members of generations g_1 , g_4 , g_5 , g_8 , and g_9 have a lower level of well-being (W_{low}) than the members of generations g_2 , g_3 , g_6 , g_7 , and g_{10} (who have a higher level of well-being, W_{high}). Since levels of well-being alternate but start with W_{low} , PE holds that there are complaints held by the members of four generations: g_4 , g_5 , g_8 , and g_9 . According to PE&BOP, the




Figure 4.4: Cosine-step




Figure 4.5: Level




Figure 4.6: Up-step

members of those generations have a complaint given by:

$$c_{i,j} = c((W_{\text{high}} - W_{\text{low}})) \quad (4.18)$$

In Sine-step, there are $4 \cdot n$ such complaints.

Cosine-step (Figure 4.4) is the inverted sequence of Sine-step. Here, the members of generations g_2 , g_3 , g_6 , g_7 , and g_{10} have the lower level of well-being, W_{low} , and the members of generations g_1 , g_4 , g_5 , g_8 , and g_9 , have the higher level of well-being, W_{high} . Since levels of well-being also alternate, but begin with W_{high} , PE holds that there are five generations whose members have complaints: g_2 , g_3 , g_6 , g_7 , and g_{10} . According to PE&BOP, the members of these generations likewise have complaints given by (4.18). Thus, in Cosine-step, there are $5 \cdot n$ such complaints — more than in Sine-step. There is greater disvalue from inequality in Cosine-step than in Sine-step, and Sine-step is thus better than Cosine-step (both with respect to inequality, and all things considered since the total well-being is the same). This appears plausible, though intuitions regarding the two cases might not be as strong as in the Simple Temporal Case.

Consider the distribution depicted in Level (Figure 4.5). In Level, the members of all 10 generations have the same level of well-being, which is the average of W_{low} and W_{high} . (Level is equivalent to (b^*) .) Here, there are no complaints regarding inequality. Thus, Level is better than Sine-step and Cosine-step, both with respect to inequality and all things considered (since the total well-being is the same). This again seems plausible. There is not a steady increase in Sine-step — some generations are worse off than prior generations. It does not seem that we should have the same indifference between Sine-step and Level as between (a^*) and (b^*) .

This result will be rejected by some. They will argue that (a^*) is *not* as good as Sine-step because there are *more* people at even lower levels of well-being. If (a^*) is not as good as Sine-step, but is no worse than Level or (b^*) , then this would seem to generate a contradiction: Sine-step is not worse than (a^*) , (a^*) is not worse than Level, but Sine-step — as just noted — is worse than Level. But we may set this problem aside. The reason for the objection, as presented here, could be derived from either maximin or prioritarian concerns. If the objection stems from maximin concerns, it would seem, then, that to combine a maximin principle of inequality with a scope restriction of complaints may not be plausible, but this does not create a problem for other views of inequality which do not support maximin. If the objection stems from prioritarian concerns, this would also affect the ranking of (a^*) relative to Level (or (b^*)), since on these views it is not true that (a^*) is no worse than Level. As a result, this view would not generate the non-transitivity when combined with the scope restriction of complaints. Prioritarianism, however, is not a *comparative* theory, and so I have set it aside. For now, let us continue to focus on the problem for individualistic egalitarianism of deciding between TNE and PE.

Finally, consider Up-step (Figure 4.6). In Up-step, the members of the first five generations have well-being W_{low} , and the members of the later five generations

have well-being W_{high} . As a result, there are no complaints regarding inequality in Up-step. Up-step is thus equally as good as Level and better than both Cosine and Sine. Is this plausible? At this juncture, a similar prioritarian worry might creep back in. With a less impressive form of progress, we might be less inclined to think that Up-step is as good as Level. The question is, why? I set this question aside for now. What I have shown, however, is that there is a form of temporal bias that seems compatible with egalitarianism — at least *prima facie*. It does not directly contradict the egalitarian insight that it is bad for some to be worse off than others through no fault of their own (as Progressive Welfarism did) but it also does not apply the insight equally across time. Depending on our intuitions, this may seem more or less plausible. And as we will see, PE needs a fair bit of refinement.

However, if we find this sequence of cases troubling, and are convinced by the worry just noted — that Sine-step is not worse than (a*) and that Up-step is perhaps not as good as Level — then we might be tempted to conclude that our initial reactions to the Simple Temporal Case were misguided. Egalitarians should declare (b*) the clear winner, and accept TNE.

Unfortunately, the answer is not quite so simple. Time wreaks far more havoc than we've considered thus far. Let us turn to a more complicated case.

4.4 THE COMPLEX TEMPORAL CASES & THE PIGOU-DALTON CONDITION

One way to compare different inequalities is by considering whether a direct transfer of the object of comparison from one person to another improves or worsens a distribution with respect to inequality. Such transfers are called *even* transfers when nothing is lost: the amount lost by the transferor is the same as the amount gained by the transferee. By contrast, *efficient* transfers are transfers where more is

gained by the transferee than is lost by the transferor, and *inefficient* transfers are transfers where less is gained by the transferee than is lost by the transferor.⁴²

Suppose that direct transfers of well-being from one individual to another are possible.⁴³ It has been argued that any measure of inequality must have the minimal property that transfers from the worse-off to the better-off increase the inequality. Sen termed this condition the “Pigou-Dalton Condition” in honor of those who first made note of it.⁴⁴ Insofar as inequality has disvalue, it would seem that this condition should also apply to the normative evaluation of inequality. That is, the Pigou-Dalton Condition also governs the betterness (regarding inequality) of transfers.⁴⁵ The (normative) Pigou-Dalton Condition states that even transfers from the worse-off to the better-off make a distribution worse. By extension, even transfers from the better-off to the worse-off “that reduce or remove the gap between them” improve a distribution.⁴⁶ Note that whether actually making these transfers is permissible is a separate question.

To illustrate, consider the following distribution between four contemporaneous people:

⁴² I here follow Temkin’s terminology. (Temkin, 1993)

⁴³ Of course, direct transfers of well-being are not possible. But direct transfers of resources (e.g. money) are possible. Hence, well-being is (in a sense) transferable: “transfers of sources of welfare will be correlated to varying degrees with changes in welfare” (Temkin, 1993, 77). For our purposes — discussing the relative value of various well-being distributions — the simplifying assumption that well-being is directly transferable will be useful.

⁴⁴ Sen (1973, 27). He writes that Dalton argued for this condition as early as 1920, but in doing so was following the lead of Pigou (Dalton (1920, 351); Pigou (1912, 24)).

⁴⁵ While this condition is commonly endorsed by economists in discussing economic inequality (Temkin, 1993, 71), it appears to have received little attention in the philosophical debate, at least under this name (e.g. neither “Pigou” nor “Dalton” appear in Clayton and Williams (2000)). For discussion, see Temkin (1993, 71f).

⁴⁶ This is Temkin’s qualification. For a discussion of why it is needed, see Temkin, 1993, 82-87.

One might think that this condition on the betterness of distributions regarding inequality is obvious: the transfers noted, from better-off to worse-off, decrease a distribution’s inequality descriptively-speaking. If inequality has disvalue, then decreasing it descriptively seems like it should also decrease the disvalue of it (normatively). However, not all descriptive measures of inequality fit this mold. This may be because of the frequent confusion between descriptive and normative measures of inequality. For a discussion of this confusion and distinction, see Sen, 1973. Even if it is, this should make it all the more clear that even though the Pigou-Dalton Condition might accurately describe the conditions under which inequality increases or decreases descriptively, it is not necessarily a normative condition on how inequality affects a distribution.

$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
10	7	5	3

According to Pigou-Dalton, an even transfer of 3 units of well-being from $p_{i,1}$ to $p_{i,4}$ would improve the state of affairs:

$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
7	7	5	6

This transfer reduces inequality in the sense discussed earlier: $p_{i,2}$ no longer has a complaint against $p_{i,1}$ and $p_{i,4}$'s complaint has been reduced. On some theories, $p_{i,3}$'s complaint has also been reduced (this is true on BOP, but not on AVE).

The condition that a distribution is improved (with respect to inequality) by even transfers from the better-off to the less-well-off (up to the point where gaps have been removed) seems plausible as a condition on valuing different inequalities. Indeed, the (static) egalitarian principles discussed above satisfy the condition.⁴⁷

The Simple Temporal Case revealed that there may be a disconnect in our moral thinking about synchronic and diachronic distributions. Analyzing a more complex case in relation to the Pigou-Dalton condition may reveal another. Call this the *Complex Temporal Case*.

We proceed with all our previous assumptions except one: there may be inequality within a generation.⁴⁸ To represent these distributions, I will now employ the use of tables. Time moves forward down the vertical axis.

Consider the following distribution, (g^*):

⁴⁷ For a discussion of how a range of telic egalitarian principles satisfy the condition, see Temkin, 1993, 71-76. PE does not, though it does with respect to inequalities which it takes to matter.

⁴⁸ The assumptions held fixed are: Each generation, g_i , has the same population, n . The vector of people in g_i is $\langle p_{i,1}, p_{i,2}, p_{i,3}, \dots, p_{i,n} \rangle$. The people within a generation are all born at the same time, live for 70 years, and die at the same time. Immediately following, the next generation repeats the cycle. In such a way, the individuals within a generation have complete temporal overlap, while generations do not overlap. All individuals are the same in every morally relevant respect (equally deserving, hardworking, virtuous, etc.); the only differences are the time at which they are located and their well-being.

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	3	3	3	3
g_2	6	6	6	6
g_3	10	10	10	10

Figure 4.7: (g^*)

This parallels the distribution presented in (a^*). There is equality of well-being within a generation, but the level of well-being increases as time progresses.

Consider now distribution (h^*), where one unit of well-being has been transferred from $p_{3,4}$ to $p_{1,1}$:

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	4	3	3	3
g_2	6	6	6	6
g_3	10	10	10	9

Figure 4.8: (h^*)

According to the Pigou-Dalton condition, this transfer — an even one from better-off to worse-off — improves the outcome. That is, Pigou-Dalton — and principles of equality that satisfy it — prefers (h^*) to (g^*). According to these theories, (h^*) is better than (g^*).

This might seem counterintuitive. The transfer resulting in (h^*) has resulted in two local inequalities: one within generation g_1 and one within generation g_3 . This may be brought out more sharply by considering a further distribution, (i^*), where the same transfer has been made again:

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	5	3	3	3
g_2	6	6	6	6
g_3	10	10	10	8

Figure 4.9: (i^*)

The Pigou-Dalton condition likewise approves this transfer. That is to say, according to theories that satisfy the Pigou-Dalton Condition, (i^*) is better than (h^*) is better

than (g^*). But if the first result was counterintuitive, so too is this: the inequality in (i^*) seems even worse than in (h^*) since the two local inequalities are even worse. Even if we had accepted TNE, the temporal egalitarian theory that applies telic egalitarian principles across time without alteration, and concluded that an even distribution over time (everyone has $w_{i,j} = 5$) is better than an uneven distribution, this method of valuing more complex inequalities over time falls short. There is strong reason to doubt a straightforward application of the Pigou-Dalton condition over time.

Why is this? While the Simple Temporal Case revealed that there may be a disconnect in our moral thinking about synchronic and diachronic distributions, the Complex Temporal Case reveals something else: namely, that we seem to care about inequalities *within* a generation in addition to inequalities between individuals across history as a whole.⁴⁹

To see this, compare the two following distributions:

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	5	5	5	5
g_2	10	10	10	10
g_3	15	15	15	15
g_4	20	20	20	20

Figure 4.10: Increasing Progression

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	20	15	10	5
g_2	20	15	10	5
g_3	20	15	10	5
g_4	20	15	10	5

Figure 4.11: Repeating Progression

According to TNE, these two distributions are equally good: they have the same historic total of well-being, and contain the same complaints regarding inequality.

⁴⁹ Note: we are still assuming non-overlapping generations. Overlapping generations raise further complications.

The only difference between the two is how in time the inequality is arranged, which, according to TNE, does not matter.

But intuitively, these are two very different distributions. In the Increasing Progression, there is complete equality at a generation, and increasing levels of well-being between generations as time progresses. We might describe this as a series of (b)-type distributions at a generation resulting in an (a*)-type distribution over time. The Repeating Progression, by contrast, is series of unequal distributions at a generation that repeats in each generation. This might thereby be described as consisting of (a)/(c)-type distributions at a time (within a generation), repeating so as to create at (b*)-type distribution over time.⁵⁰ As argued in the last section, the value difference between an (a*)-type distribution and a (b*)-type distribution is (at least *prima facie*) different than the value difference between an (a)/(c)-type distribution and a (b)-type distribution. The Repeating Progression is *prima facie* worse than the Increasing Progression.

One way to explain why is to see that in the Repeating Progression there is an inequality that is *preserved* over time: each generation contains the same synchronic inequality as the generation before it. Since this synchronic inequality does not change as time progresses, we might say that it is preserved; there is no improvement in the state of affairs over time with respect to inequality. (There is also no improvement with respect to well-being in the Repeating Progression, unlike in the Increasing Progression.) If inequality is the sort of thing which may be improved upon — indeed, it is not uncommon to at least speak as though it is⁵¹ — then this might explain a preference for the Increasing Progression over the Repeating Progression. And the contrast on this point is stark: while it seems to make sense to say that inequality increases or decreases over time, it does not make sense to say

⁵⁰ I say (b*)-*type* because this is not the exact equivalent of (b*) since there is inequality within generations. In (b*), all people were equally well-off, both as considered within and across generations.

⁵¹ To wit, the multitude of news stories about increasing income inequality in America.

that it increases or decreases across space. To say that an inequality over time — a *temporal inequality* — is preserved across space is bizarre. Even assuming loose connections (e.g. familial, societal, civilization) between, for example, $p_{1,j}$ (for all j), and between all $p_{2,j}$ (for all j) of the Increasing Progression, it does not seem to make sense to say that this inequality over time is preserved across space in the same way that it makes sense to say that there is an inequality across space which is preserved over time (especially if loose cross-temporal connections like family or societal relations exist between, for example, all $p_{i,1}$ (for all i) and between all $p_{i,2}$ (for all i)). There seems to be a sort of moral relationship between people who exist at a time that has no analog between people over time.⁵²

A practical example, drawn from Temkin, might help elucidate this concern:

Mr. Gelt accumulated a large fortune. The Gelt children and grandchildren have been well cared for, and are better off than most, though they have not been as successful as Mr. Gelt in amassing their own private fortunes. To avoid estate tax, Mr. Gelt begins distributing his vast fortune to his children and grandchildren.

Temkin observes: “Since, by hypothesis, each of these transfers goes from a rich to a poorer person, the egalitarian should approve these transfers *if* she is concerned about equality of *wealth*.”⁵³ But does this improve the state of affairs with respect to inequality? Temkin argues that it does not, but on the grounds that, owing to diminishing marginal returns of wealth, the even transfer of wealth has led to an efficient transfer of welfare which benefits those already better off than most, thereby exacerbating an existing inequality of welfare — the particular type of inequality that Temkin thinks matters. But it seems that there might be a different lesson to be drawn from this case.

⁵²cf. Arrhenius and Bykvist (1995, 67, 112ff.)

⁵³ Temkin (1993, 84-85).

There is another reason why “many would be suspicious of the claim that, as Rockefeller began disbursing his millions to his children, this improved inequality.”⁵⁴: such a transfer seems to further entrench a synchronic inequality. It is not merely that Rockefeller’s legacy worsened inequality of welfare — his legacy *preserves* it by directly causing the inequality of his time to persist over multiple times. The Pigou-Dalton condition, it appears, also needs revision for application to transfers over time.

To begin unpacking how to do this, consider a simplified version of even transfers of well-being called the *Inheritance Case*. The original distribution, over four generations, is as follows:

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	50	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10

Figure 4.12: Inheritance Case: Original Distribution

One of two even transfers of 30 units can be made: $p_{1,1}$, shown in blue, may either pass 30 units on to $p_{2,1}$, $p_{3,1}$, and $p_{4,1}$ (light blue) with each receiving 10, or to $p_{1,2}$, $p_{1,3}$, and $p_{1,4}$ (light green), with each receiving 10. Thus, we have a choice between the transfer leading to (A) (Figure 4.13) or the transfer leading to (B) (Figure 4.14).

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	20	20	20	20
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10

Figure 4.13: Inheritance Case: (A)

⁵⁴ Temkin (1993, 85).

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	20	10	10	10
g_2	20	10	10	10
g_3	20	10	10	10
g_4	20	10	10	10

Figure 4.14: Inheritance Case: (B)

Which yields the better distribution with respect to inequality?

To say that the inequalities in (A) and (B) are equally bad seems absurd. The transfer leading to (A) eliminates the inequality in g_1 while leaving unaffected all other generations. By comparison, the transfer leading to (B) only lessens the inequality in g_1 (by reducing the size of complaints) while introducing inequality into the subsequent three generations. This gives reason to suggest that (A) is better than (B): an inequality has been resolved without perpetuating inequalities — generational equality is preserved.

On the other hand, one might argue that there is a reason to prefer (B) to (A). In (B), well-being remains constant over time, with each generation being neither better nor worse off than the other generations. In (A), however, the first generation is better off than all the rest; well-being declines. (A) is thus a (c*)-type distribution. (B), where the level of well-being remains the same from one generation to the next, is a (b*)-type distribution. This might appear to be a reason to prefer (B) to (A). As discussed in the last section⁵⁵, (c*) is a worse distribution than (b*). Does the same apply to comparisons between (c*)-type and (b*)-type distributions? If so, as seems to be the case, there is also a *pro tanto* reason to favor (B) to (A). The question, then, is which is stronger? I will return to this in Part III.

For now, it appears that this line of reasoning about transfers parallels our discussion of the Increasing Progression and the Repeating Progression. Those cases suggested what the Inheritance Cases seem to confirm: we care about preserving gen-

⁵⁵ Above, page 37.

erational equality through time, but do not care about preserving temporal equality through space. The next question is whether there is a view that is both tenable, and able to accommodate this conclusion. But first, an objection.

4.4.1 AN OBJECTION: NO DIFFERENCE BETWEEN DIACHRONIC INEQUALITIES AND GENERATIONAL INEQUALITIES

Earlier, in our discussion of the Simple Temporal Case, we entertained the possibility that there is no difference between diachronic inequalities and generational inequalities, conceding that it might be a strong enough reason to doubt our *prima facie* intuitions about the Simple Temporal Case. But we put off making a decision. However, our discussion of the Complex Temporal Case in relation to the Pigou-Dalton condition has cast even stronger doubt on this complete temporal indifference. Before considering egalitarian positions that may be able to accommodate the Complex cases, we must meet head-on an objection that was raised earlier, and that we ceded for the moment. That is, we must meet the objection that we do not have reason to care about generational inequality to a greater extent than inequalities between all individuals over time.

Underlying this objection are two features: the fundamental insight, and an assumption about time.⁵⁶ We already accepted that as egalitarians, our proper concern is the fundamental insight — that it is bad for individuals to be worse-off than others through no fault of their own.⁵⁷ The question then is about the scope of comparison. The objector will argue that the scope of this comparison should not be limited to a comparison within generations or, more generally, that this comparison should not be limited in scope to particular times. The assumption

⁵⁶ Another route to this objection is the argument from overlapping generations, which we will discuss briefly (below, page 230). This argument cannot be used here given our assumption of non-overlapping generations.

⁵⁷ Whether we should, in fact, be egalitarians is a separate question not addressed in this thesis.

made here is that relative temporal location is morally irrelevant, and thus whether people are contemporaneous is likewise morally irrelevant. An obvious reason why is that people do not have control over the time at which they are located (even less so than over the spatial coordinates they find themselves at!), and therefore time is not a difference that an individual can or should be faulted for.

This objection might be further strengthened by an underlying metaphysical view of the nature of time, in particular the sort of four-dimensionalism that entails a strict analogy between space and time. On such a view, space and time make no difference: roughly, two contemporaneous people at different spatial locations are the same as two people in the same spatial location but at different times. Indeed, that this is a sloppy formulation of this view is revealing. The better way of formulating it is to say this: Make the assumption of point-size people, and describe points (locations) within the four-dimensional space-time continuum as (x, y, z, t) where x , y , and z correspond to the axes normally associated with space and where t corresponds to the axis associated with “time”. This critic of valuing generational equality would then argue that there is no difference for purposes of comparison between two people located at $(1, 0, 0, 0)$ and $(2, 0, 0, 0)$ respectively, and two people located at $(0, 0, 0, 1)$ and $(0, 0, 0, 2)$. Both are merely cases of being located at *different coordinates* within the four-dimensional space-time continuum. The case against valuing generational (or contemporaneous) equality could, perhaps, be further bolstered by one of the key arguments in favor of a four-dimensional picture, the argument from special relativity, according to which there are no favored planes of simultaneity: being contemporaneous is, in a deep sense, arbitrary.⁵⁸

I set the metaphysical version of the objection aside. It is not clear that it would truly bolster the case for TNE. An egalitarian who restricts scope accord-

⁵⁸ Sider (2001, esp. 42-52). It might be questioned whether we’re at least close enough to *seem* simultaneous, so that this objection would not cut as deep. The metaphysical objection raises many interesting questions that might have normative implications; however, as I will explain next, I must set them aside for now.

ing to location in time could do so according to distance from coordinates. That is, an egalitarian who restricts scope according to location in time might take the metaphysical view as reason to have similar scope restrictions in space. More importantly, there is much to be worried about from the non-metaphysical version of the objection. While the question of normative implications of different metaphysical views is an important one, wandering into hotly contested metaphysical territory will carry us too far afield. The metaphysical problem must wait.

The objection that if the fundamental insight applies, it applies equally across time, can stand without a metaphysical backing. To argue that equality matters *intrinsically* seems to suggest that we should not discriminate according to time. Consider two individuals, Mary and Harold. If it matters *intrinsically* that Mary is worse off than Harold through no fault of her own, shouldn't this matter regardless of where they are located in time relative to each other?

Reconsider the argument from displacement of persons against discounting over time presented earlier in our discussion of TVV*.⁵⁹ Mary was an ordinary human being living in 2009 with a well-being of 10. She lived in a (b*) world: everyone who ever existed or will exist in the history of her world also had a well-being of 10. I argued that if we were somehow able to move Mary in time without affecting anything else, that it would not affect the total goodness of the historic distribution: the world in which Mary is located in 2109 is no better — and no worse — than the world in which Mary is located at 2009.

This argument seems to bolster the idea that the location of a person (and, consequently, their well-being) in time should not affect the goodness of the historic distribution. There is a key difference, however, between the case of Mary and the cases we have been considering: Mary's world is a (b*) world (not just a (b*)-*type* world).

⁵⁹ Above, page 42.

What if Mary lived in an Increasing Progression world? That is, what if Mary lived in an (a*) world? Mary still actually lives in 2009 with a well-being of 10. People located earlier in time enjoy less well-being, and those located later in time enjoy higher well-being. Suppose, again, that we are able to move Mary in time without affecting anything else. How does this affect the goodness of the historic distribution?

Suppose those living in 2109 have well-being of level 20, and that those living in 1909 have well-being level 5. Moving Mary to either of these will create a local inequality at those times that had not been there without the move. These worlds — one where Mary has been moved to 2109 and one where Mary has been moved to 1909 — appear to be no better than the world where Mary actually lives. But are they any worse?

There are three possible answers:

1. The worlds are no worse. All worlds are equally good.
2. Both worlds are worse owing to the local inequalities created by moving Mary.
3. The world where Mary is moved to 2109 is worse since both a local inequality is created and inequality worsens over time. The world where Mary is moved to 1909, however, is less bad — and perhaps no worse than the actual world — since inequality becomes better over time.

The third may be plausible, but I set it aside as a further complexity. The first option does not seem particularly plausible: there is a difference, namely that local inequalities are created, and inequality is bad. Even though Mary is no worse off (or better off) relative to the historic population, she is relative to her contemporaries. This, as I have been arguing, seems intuitively problematic. It seems plausible to think that it at least provides *some* reason to prefer the original world — to answer with the second available response.

We seem to have arrived at an impasse. But there is a final point to raise before moving on that strikes at the two assumptions outlined at the beginning of this section. The fundamental egalitarian insight advocated earlier states:

It is bad for one person to be worse off than another through no fault of their own.

This seems plausible, but perhaps it is not quite correct. Underlying the fundamental egalitarian insight seems to be another notion of fairness:

It is bad for a distribution to “fail to reflect the presence or absence of morally relevant differences between [the individuals in the distribution].”⁶⁰

On the fundamental egalitarian insight, fault seems to be taken to be the only morally relevant difference. In cases where we must make choices between different inequalities, might time be another?

4.4.2 THE COMPLEX TEMPORAL CASE: IMPLICATIONS FOR PE & TNE

Given that our goal (in Part I) is not to defend a particular position, but to examine the implications of applying — or not applying — these principles over time, we must consider what the implications of the Complex Temporal Cases are for the principles presented earlier in our discussion of the Simple Temporal Case, TNE and PE. Table 4.1 provides the relative rankings of distributions provided by each of these views in response to the cases presented.

⁶⁰ Kagan (1998, 54).

Table 4.1: TNE and PE Evaluations of Cases

Egalitarian View	Simple Cases			Complex Cases		
	Simple Case	Simple Temporal Case	Temporal Leveling Down Revisited	Complex Temporal Case	Increasing vs. Repeating	Inheritance Case
(E)	(b) (a), (c)	—	—	—	—	—
TNE & BOP	(b) (a), (c)	(b*) (a*), (c*)	(f*) (d*) vs. (e*) depends on weight given to complaints	Depends on Weighting <u>AP</u> Tie <u>WAP</u> (i*) (g*) (h*)	Tie	Tie
<i>Key:</i>	(E) TNE PE BOP AVE AP WAP GE	Pure Individualistic Egalitarianism Temporally Neutral Egalitarianism Progressive Egalitarianism Complaints Relative to Best-off Person Complaints Relative to Average Additive Principle of Equality Weighted Additive Principle of Equality Generational Equality				

Table 4.1: TNE and PE Evaluations of Cases (continued)

Egalitarian View	Simple Cases			Complex Cases		
	Simple Case	Simple Temporal Case	Temporal Leveling Down Revisited	Complex Temporal Case	Increasing vs. Repeating	Inheritance Case
TNE & AVE	(b) (a), (c)	(b*) (a*), (c*)	(f*) (d*) vs. (e*) depends on weight given to complaints	(i*) (g*) (h*)	Tie	Tie
PE & BOP	(b) (a), (c)	(a*), (b*) (c*)	(d*), (f*) (e*)	(g*) (h*) (i*)	Increasing Repeating	Tie
PE & AVE	(b) (a), (c)	(a*), (b*) (c*)	(d*), (f*) (e*)	(g*) (h*) (i*)	Increasing Repeating	(B) (A)
GE only (static)	(b) (a), (c)	Tie	Tie	(g*) (h*) (i*)	Increasing Repeating	(A) (B)
<i>Key:</i>	(E) TNE PE BOP AVE AP WAP GE	Pure Individualistic Egalitarianism Temporally Neutral Egalitarianism Progressive Egalitarianism Complaints Relative to Best-off Person Complaints Relative to Average Additive Principle of Equality Weighted Additive Principle of Equality Generational Equality				

We already saw that the Pigou-Dalton condition yields counterintuitive results when evaluating inequalities over time. Unsurprisingly, TNE likewise yields the counterintuitive results. TNE approves the cross-temporal Pigou-Dalton transfers in the Complex Temporal Case that appeared problematic. Moreover, it fails to distinguish between the outcomes in any of the other Complex Cases presented (Increasing and Repeating Progressions, and the Inheritance Case). In the New Displacement of Persons Case, TNE would have to provide the answer that we argued was implausible: not only does displacing Mary in time in an (a^*) world not affect the goodness of the historic distribution on TNE, displacing Mary cannot even be said to worsen the inequality. These results, coupled with the concerns sketched out in the discussion of the Simple Temporal Case, strongly suggest looking for an alternative to TNE.

The alternative suggested in our discussion of the Simple Temporal Case was Progressive Egalitarianism (PE). But PE encounters problems when confronted with the Complex Cases. PE yields plausible results in response to the Complex Temporal Case, and is able to distinguish between the Increasing and Repeating Progressions. But when combined with BOP and AVE, PE offered counterintuitive responses to the Inheritance Cases. In the Temporal Displacement Case, PE is able to provide the intuitive result that displacing Mary affects the outcome, but how it weights the two alternatives is not straightforward: it depends on which view about what complaints are relative to is assumed, and on how many individuals there are in a given generation. Since this particular problem is only dealing with a single generational inequality, I set it aside for a moment; it may not turn out to be a problem.

The bizarre response to the Inheritance Cases, however, leads to several questions: What is the source of the bizarre inheritance result? Is it truly problematic? And can it be resolved?

The source of the bizarre inheritance result is that PE treats complaints against past and present the same.⁶¹ In the inheritance cases, the better-off are concentrated in the past in (A) rather than being spread evenly through time as they are in (B). The result is that, when PE is combined with BOP, the inequality is measured to be the same because the complaints created by local, generational inequalities in (B) are the same as the complaints created by the inequality between the first generation and the others in (A). When combining PE with AVE, the fact that the better-off in (A) are concentrated in the past instead of spread out over time becomes significant: the average against which the worse-off's complaints are measured is higher in (A) than in (B) respectively. As a result, PE & AVE ranks (B) as preferable to (A).

Is this a problem? Suppose we add more generations to each (A) and (B). If the generations added are all equal, with each member of the new generation having well-being 10, then the complaints derived from the individuals of each new generation are the same for both (A) and (B), since in both all the better-off are located in the past so that the average of past and present welfare, by which complaints are measured, is now the same. There will still be a preference for (B), but while the difference between the two will remain the same, this difference will become smaller in comparison to the total complaints. The judgment then, is that a temporary series of generational inequalities is less bad than the same number of better-off being located all at once in the past: for a limited number of better-off, it is better to spread them over time. But this looks almost exactly like the counterintuitive result we rejected when we argued that sustained inequalities within generations (repeating generation inequalities) is worse than one generation being better-off. Should it matter that the better off generation is located in the past?

⁶¹ To use an analogy with the nature-of-time debate, PE treats all inequalities in the past and present as equally real.

Perhaps, it does make a difference. Suppose, for comparison, that the generations added to (B) continue to repeat the inequalities ($w_{i,1} = 20$ and $w_{i,2} = w_{i,3} = w_{i,4} = 10$) while the generations added to (A) are as before equal but lower than the first generation ($w_{i,1} = w_{i,2} = w_{i,3} = w_{i,4} = 10$). The size of additional complaints in (B) will remain constant, while the size of additional complaints in (A) will continue to shrink as the average does. Eventually, a point will be reached where the inequality in (B) is worse than the inequality in (A). But while this may sound plausible — perhaps this is a justified, though unexpected, result — this use of PE & AVE has its own problems. The problem is that PE & AVE will still fail to distinguish between an inequality with someone *very much* better-off in the distant past as compared with a sizeable inequality relative to those in the closer past or present. This, indeed, is more a problem with AVE than with PE, and parallel problems arise in static cases. We will not discuss it here; that it fails to distinguish the two scenarios just described is reason enough to see the combination of AVE with PE creates more of a problem than it solves.⁶²

Is the result — a tie — that is yielded by PE & BOP a difficulty? Perhaps not: there are *pro tanto* reasons in favor of (or against) each one. (A) is a (c*)-type distribution, which, as argued in the Simple Temporal Case, is less favorable than a (b*)-type distribution like (B). By contrast, (B) contains generational inequalities that (A) does not, which constitutes a *pro tanto* reason against (B) and in favor of (A). Depending on how the relevance of such inequalities are weighted, the *pro tanto* reasons might balance each other out, yielding a tie. This may seem plausible, but it doesn't seem particularly implausible given that (A) is a (c*)-type — and not an (a*)-type — distribution.

Adding generations, however, yields a problem. Consider two distributions, (A') and (B'), depicted in Figures 4.15 and 4.16, where the distributions of (A) and (B)

⁶² For an extended discussion of why we should reject AVE, see Temkin, 1993.

have been set within matching extended histories. (The better-off are highlighted in light blue for the reader's convenience.) PE & BOP also ranks these two dis-

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	10	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	20	20	20	20
g_9	10	10	10	10
g_{10}	10	10	10	10
g_{11}	10	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	10	10	10	10
g_{18}	10	10	10	10
g_{19}	10	10	10	10
g_{20}	10	10	10	10

Figure 4.15: (A')

tributions the same. (A') seems intuitively more plausible: the temporal inequality is an anomaly, and there are no generational inequalities. By contrast, (B') has a repeating generational inequality, and temporal inequalities (all those in g_9 or later have a complaint regarding inequality relative to the past). But, perhaps, this is not enough to convince that PE is in need of revision: one might argue that given that the series of repeating inequalities in (B') is short-lived, (B') is not particularly worse than (A') after-all.

Consider one final distribution comparison, (A'') and (B'') (Figures 4.17 and 4.18). The problem, I suggested, is that PE treats complaints relative to the

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	10	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	20	10	10	10
g_9	20	10	10	10
g_{10}	20	10	10	10
g_{11}	20	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	10	10	10	10
g_{18}	10	10	10	10
g_{19}	10	10	10	10
g_{20}	10	10	10	10

Figure 4.16: (B')

better-off the same so long as the better-off are located either earlier than or contemporaneously with the individual with the complaint. These last two distributions make this clear. According to PE&BOP, the complaint of $p_{20,4}$ (shown in green) regarding inequality is the same in both (A'') and (B''). This is because complaints relative to better-off in the *very* distant past are weighted the same as complaints relative to the better-off in the same generation. This is problematic. For long histories, PE&BOP would deliver the judgment that a distribution with a very well-off early generation is worse than a distribution with gross inequalities in later generations but no earlier problems because of the cumulative effect of complaints relative to the distant past. In such a way, PE approves Pigou-Dalton transfers from earlier to later, even when these result in more sustained inequalities at later generations. PE falls into a similar trap as TNE.

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	20	20	20	20
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	10	10	10	10
g_9	10	10	10	10
g_{10}	10	10	10	10
g_{11}	10	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	10	10	10	10
g_{18}	10	10	10	10
g_{19}	10	10	10	10
g_{20}	10	10	10	10

Figure 4.17: (A'')

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	10	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	10	10	10	10
g_9	10	10	10	10
g_{10}	10	10	10	10
g_{11}	10	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	20	10	10	10
g_{18}	20	10	10	10
g_{19}	20	10	10	10
g_{20}	20	10	10	10

Figure 4.18: (B'')

Can this problem be resolved? PE emerged as a result of questioning the scope of comparative complaints. If there is reason to restrict these to certain times, there might also be reason to discount complaints according to the temporal distance between the two individuals being compared. The greater the temporal distance between the better-off and the worse-off, the less the complaint is weighted. Call this the *Temporal Comparison Thesis*:

Temporal Comparison Thesis (TCT) An individual's complaints against those who are temporally distant will have less weight than his complaints against those who are temporally close or contemporaneous.

Call this theory, which combines the TCT with the Progressive Scope Restriction of PE, *Weighted Progressive Egalitarianism* (WPE). The details of such a weighting need to be worked out — a task that must wait until later in the thesis — but doing so appears to be worthwhile. Even before such details are determined, it is clear that WPE has certain advantages. First, it would help avoid some of the more egregious problems above. Second, it would also serve to explain why generational inequalities are more grave than temporally extended inequalities but without having to excessively adjust our formulas to include both a component for individual complaints regarding generational inequalities and a component for individual complaints regarding temporally extended inequalities. Third, it avoids a problem regarding groups that might be raised given our focus on generational inequalities. Fourth, it allows a rebuttal to the argument for TNE from overlapping individuals, thereby accommodating both an egalitarian position and our temporal intuitions. Finally, it seems poised to shed light on a problem in population axiology. I will address the first three in this section. The other implications I will discuss in the final chapter.⁶³

⁶³ Another case where WPE yields plausible results is the *Resetting Civilization Case*: those immediately after the fall of a civilization have strong complaints regarding inequality; those worse off who are temporally distant — and properly considered part of another civilization — have

First, WPE would help address the more egregious problems that PE on its own cannot. By discounting complaints relative to temporally distant individuals, WPE can make the distinction needed in comparing (A'') and (B''): $p_{20,4}$'s complaint is *not* the same in both distributions — it is much stronger in (B''). In addressing the original Inheritance Case and the Extended Inheritance Case, WPE could be adjusted so as to decide in favor of making even transfers of well-being to members of the same generation — or temporally close generations — without being committed to indifference between making even transfers to the same or temporally close generations and generations very far in the future.

Second, it is able to achieve this without positing a separate “communal” good of equality at a generation. That is, WPE allows complaints against contemporaneous inequalities to take precedence without separating the way in which we (dis)value generational inequality and temporal inequality. It falls out that generational inequalities are worse than temporally distant inequalities not because equality-at-a-time matters *per se*, but because temporally close inequalities matter more than temporally distant ones. This, coincidentally, is why I set the metaphysical debate aside: a B-theoretic four-dimensional space-time continuum theory of time does not necessarily provide reason to be temporally neutral egalitarians. Rather, it might provide reason to weight “spatially” distant inequalities as well.

Third, it is able to achieve this without appeal to concern about equality between groups. This is important. Temkin rightly notes:

After all, if it is bad for one person to be worse off than another through no fault of her own, this should be so whether or not they happen to be lumped together as members of the same group.⁶⁴

weaker complaints against the prosperity of the past. This raises an interesting point I hope to explore further in another work: a parallel between the morally relevant R-relation connections of personal identity and those connecting societies or moral communities.

⁶⁴ Temkin (1993, 101).

What matters is inequality between individuals. Indeed, Temkin makes a sustained argument against concern about inequality between groups, summarized in his words as follows:

[I] considered two questions regarding the proper focus of egalitarian concern. I began by considering the average and proportional average principles of equality, both of which focus concern on reducing inequality between the average members of society's groups. I argued that both face profound difficulties, not the least being that each gives weight to small inequalities between groups while giving *no* weight to (even) large inequalities within groups. Our discussion reinforced one of the book's central tenets. Claims expressing concern about society's inequality, or about inequality between society's different groups, must ultimately be understood as expressing concern about the inequality between the society or groups' *members*. We have no special concern about society *per se*, or about the groups *themselves*, what we have is a concern for the *individuals* who compose society and its groups. Equality, like many other ideals, should be understood individualistically, not holistically.⁶⁵

The trouble is that, as the Complicated Temporal Case reveals, the individualistic approach encounters difficulty when applied evenly over non-overlapping generations. Concern about *individuals* without concern about *localities* creates indifference to or preference for perpetuated inequalities through time. Yet, without appealing to membership *in* a generation, WPE is able avoid this difficulty facing the individualistic approach. This is *not* to say that it is not bad when a person in the distant future is worse off than a person in the distant past. Rather, it is to say that this inequality is less bad than the same inequality between contemporaneous people. WPE is simply claiming that inequalities between *individuals* are less bad

⁶⁵ Temkin (1993, 285). For the extended discussion, see Chapter 4 of the same.

when those individuals are temporally distant.

In response to the cases shown, then, WPE emerges as the most plausible option.⁶⁶ We turn now to the work of evaluating a principled basis for WPE.

⁶⁶ There might remain a worry that egalitarianism violates the Independence of the Utilities of the Dead Axiom:

Independence of the Utilities of the Dead “the ranking of any two alternatives [is] independent of the existence of individuals whose lives are over and who had the same lifetime utilities, birth dates, and lifetimes in both alternatives.” Blackorby et al. (2005, 253)

This axiom has also been called “Independence of the Existence of the Dead”. For a discussion of its implications, see Blackorby et al. (2005, 253); also, Blackorby et al. (1995, 1311ff.).

To determine whether a particular version of egalitarianism violates this axiom, more information is needed about how complaints are calculated and weighted. I suspect that, for fixed number cases, the utilities of the common dead of two alternatives will not change the rankings of those two alternatives on egalitarian views which reject weak leveling down (the view that leveling down might not make an outcome worse; that is, the condition on egalitarianism that the all-things-considered value function must be increasing in total well-being). But I will not pursue the matter here.

Even supposing that egalitarian views violate Independence of the Utilities of the Dead, I do not think this a reason to reject egalitarianism. Suppose (contrary to our simplifying assumptions) that we have two individuals who have some partial overlap, Elder and Younger. Most egalitarian views assume that inequalities of lifetime well-being have at least some disvalue. *cf.* McKerlie (1989, 491). If Elder is better off than Younger with respect to lifetime well-being, then Younger will have a complaint against Elder. This seems like a plausible egalitarian position. Call this History Overlap.

Suppose now that there is an alternate history, History No-Overlap, in which Elder dies prematurely just a few years earlier but is still far better off than Younger. As a result, in History No-Overlap, Elder and Younger are never alive at the same time. But the temporal gap between them is small: we can suppose that it is a matter of a few minutes. If there was a reason to think that Younger suffered an unfairness in History Overlap, I do not see how we can say that he does not in History No-Overlap. Elder may have lived the high life, using and hoarding resources, and getting as much from those resources as Younger could have. This was unfair to Younger in History Overlap. That Elder is, say, hit by a car a few years earlier in History No-Overlap does not make this distribution any less unfair for Younger. Temporal distance might make a difference, as I have argued, but this minimal amount of temporal distance is insufficient to eliminate Younger’s complaint.

If Younger still has a complaint against Elder in History No-Overlap, then the existence of (at least some of) the dead do affect the personal good of the living. This opens the door for violation of the Independence of the Utilities of the Dead. But as History No-Overlap shows, I do not think this is a reason to reject the comparative view. History No-Overlap also does not appear to be the case which gave rise to the axiom. Rather, what appears to have motivated this axiom is “the idea that the existence of people who are *long* dead should not influence the ranking of feasible alternatives (which necessarily have a common past)” (Blackorby et al., 2005, 254, emphasis added). Indeed, Blackorby et al. acknowledge that history must matter to some extent otherwise “it is impossible to distinguish between a person dying just before period t and a person not being born in period t ” (Blackorby et al., 1995, 1311). If there is away to avoid the influence of the *long* dead without abandoning egalitarianism, then the egalitarian is not forced to accept this axiom.

Part II

The Problem of Temporal Distance

At the end of Part I, I suggested that a turn to Weighted Progressive Egalitarianism helped alleviate many of the problems posed by the temporal cases. But the question remains whether such a position can be justified. This section begins answering that question by exploring whether and how the temporal location of agents and objects matters. I argue that it does, and, importantly, that it does so in a manner which supports both the Progressive Scope Restriction and the Temporal Comparison Thesis.

The argument in favor of this conclusion is not a straightforward one. In particular, it depends on answers to questions that do not, at first, seem immediately related to the problems of the previous section about whether inequalities between temporally distant people matter, and about whether egalitarians ought to abandon temporal neutrality when valuing complaints between temporally distant individuals. Rather, this argument begins with claims.

What is a claim? A claim most frequently manifests itself as a demand or assertion for something – be it a material object (e.g. medication) or a form of treatment (e.g. respect, consideration for a job) – that is *owed* to a particular individual. But that individual need not *actually* demand that they be given what it is owed. An individual can have a claim to something even if they do not assert it, and even if they are unaware that they have one.

This is due to the structure of a claim. Implicit in claims are three features: a subject, an object, and a basis. The basis is some *reason* in virtue of which the subject should receive the object to which they have a claim. This reason is an objective, not a subjective, reason, and exists independently of the subject's awareness of it.⁶⁷ Moreover, since the claim exists in virtue of this basis, an agent

⁶⁷ An objective reason is one which depends on facts; it is independent of an individual's perspective. A subjective reason is a reason from the individual's perspective, given what they know and what they believe to be the facts. These might be better termed "true" reasons and "apparent" reasons. They may differ, as when an individual has reasons of which they are unaware. For example, if a medication were to be laced with poison, this is a reason that the agent should not take the medication (ignoring special cases where the death of the agent is the desired result).

need not assert their claim in order for the basis to exist; rather, the basis comprises the justification for an assertion of a claim.⁶⁸ In other words, claims take the form “A has a claim to X in virtue of Y” where A is the subject, an agent of the relevant sort; X is the object, a good or bad of the relevant sort (e.g. well-being, treatment, respect); and Y is the basis, some fact that provides a *pro tanto* reason for why A should receive X. There is nothing built into the structure of a claim that requires the awareness of the agent for the existence of the claim because there is nothing built into the structure of the basis that the agent must be aware of the reason or assert their claim for it to be a reason why they should get the object claimed.

Claims are important because claims are at the root of *fairness*, and the badness of inequality is that it is unfair.⁶⁹ The *Fundamental Egalitarian Insight* might then be rephrased:

It is bad for those with equal claims to receive unequally.

In such a way, claims ground the value of comparisons. They explain why differences in what people receive *matter*. I will return to the relationship between claims, fairness, and the disvalue of inequality in Part III.⁷⁰

Our question in this part, then, is whether *mere* temporal distance can affect the relative strength of claims. If it can be demonstrated that mere temporal distance does affect the relative strength of at least some claims, this might provide grounds for the Progressive Scope Restriction and Temporal Comparison Thesis. If not, then support will need to be found elsewhere, or these principles abandoned.

But the agent might be unaware that the medication is poisoned. In this case, the agent lacks a subjective reason (a reason of which they are aware) to refrain from taking the medication. But there remains a very important *objective* reason why the agent ought not to take the medication: it is laced with poison. Ignorance of reasons does not mean that they do not exist.

⁶⁸ A possible exception may be certain types of conventional claims where it is built into the agreement that the agent must assert their claim in order to retain it (e.g. some powerball lotteries).

⁶⁹ *c.f.* Broome (1989).

⁷⁰ *infra* 174.

I hope that this section will be of interest to anyone with an interest in the success or failure of telic egalitarianism. It effectively constitutes an evaluation of the strength of my argument against telic egalitarianism found in Part I. If the PSR and TCT can be defended, this constitutes a rebuttal to my arguments of Part I. But if the PSR and TCT fail, this provides strong reason to reject telic egalitarianism; I do not know how else it might be salvaged.

This section divides into three chapters. The first sets out a toy theory of claims. The second considers and rejects a principle about temporal irrelevance in dividing claims in a world with a single resource that is growing. The third analyzes how, given that time is not irrelevant to derivative claims to resources, time might be relevant.

AN OUTLINE OF A TOY THEORY OF CLAIMS

In this chapter, I set out a toy theory of claims. I do not mean to suggest that this is a fully worked-out theory of claims. Rather, I hope that the “toy theory” includes enough of the features that such a worked-out theory would need to include that it will serve the purposes of our discussion about how claims work over time. Developing a fully defensible theory of claims would take us too far afield from our present inquiry, and so is work left to another time.

5.1 WHAT IS A CLAIM?

As explained in the introduction to Part II, a claim is a three-place relation between a subject, an object, and a basis, where the basis is some reason in virtue of which the subject has a claim to the object. They take the form “A has a claim to X in virtue of Y” where A is the subject, an agent of the relevant sort; X is the object, a good or bad of the relevant sort (e.g. well-being, treatment, respect); and Y is the

basis, some fact that provides the reason for why the *individual*, A, should receive X.

Claims may have different types based on the type of reasons which support the claim, and based on the types of objects that can be so claimed.¹ Two familiar types of claims are desert claims — claims to that which is deserved — and entitlement claims — claims to that which one has been entitled by agreement (e.g. by a contract).² They differ primarily as to what counts as a basis for the claim, as will be discussed. For example, in the case of desert claims, the basis must be some fact *about the agent*. In the case of entitlement claims, the basis might extend to include some fact about an *agreement* or *contract* made between the agent and another agent or institution.

A claim is one reason that an individual should receive a given resource or good, but it is a peculiar sort of reason, differing from other reasons in this fundamental way: it is *intrinsically bad*³ if an individual's claims are not satisfied, in addition to the loss of benefit that would thereby have accrued.⁴

To illustrate, consider a simple case: David wants a boat because, he believes correctly, it will give him pleasure to own a boat. That the boat will give him pleasure is a *pro tanto* reason for David to acquire a boat. This is the straightforward welfarist reason: were David to acquire a boat, David's well-being would increase, and, since David's well-being has value, this improves the outcome (at least in this respect). But while receiving a boat would make David better off, there is nothing *bad* about a state of affairs in which David does not receive a boat above and beyond the difference in the value of well-being that he would have received.

¹ Might there be a further way of distinguishing types of claim based on which agents qualify to have such a claim? I suppose there may be, but I will not address it here.

² For an explanation of the difference between moral desert and entitlement, see Griffin (1986, 255); Kleinig (1971, 74).

³ That is, it detracts from a state of affairs. *Supra* 5.

⁴ Feldman describes desert in a similar manner: "Roughly, to say that a person deserves some good is to say that it would be 'distributionally appropriate' for him to get it" (Feldman, 1995a, 161).

This is considerably different from a case where David has a claim to the boat. For sake of argument, suppose David had entered into a raffle for a boat and won. David now has an entitlement claim to the boat in virtue of his winning ticket. There is a reason that *David* should get the boat (as opposed to anyone else who would similarly benefit from receipt of the boat). Is there anything bad — above and beyond the difference in well-being that will result — if David’s claim to the boat is not satisfied? If the satisfaction of this type of claim has value — if the satisfaction of this type of claim matters — then there is: the difference between the outcome where David receives the boat and the outcome where he does not is not only the difference in David’s well-being, but also the denial of his claim. It is not only bad that David did not experience the increase in well-being; rather, it is *also* bad that his claim was denied.

This is not to say that claims may never be overridden. They provide only a *pro tanto* reason why the individual should receive that to which they have a claim: their satisfaction has value, but this value must be weighed against other sorts of goods.⁵ So conceived, this is a theory of claims designed for a consequentialist framework.⁶ But I hope that the work presented here may be of use to proponents of other moral theories as well, as I suspect similar arguments could be made regarding the treatment and status of deontic claims.

In addition to a theory about the absolute value of the satisfaction of claims (i.e. the value of *whether* the claim is satisfied), a comparative theory is also needed since there will inevitably instances where absolute desert is violated, either because not all claims can be satisfied or some claims can be more than satisfied.⁷ In these

⁵ cf. Feldman (1995a, 161).

⁶ Arrhenius (2006) provides a useful framework for including the value of noncomparative desert within a population axiology in the line of the “fit-idea” (what has value is when the desert of an agent and what they receive matches, that is, “fits”). His suggestion is worth pursuing, but is work left to another time: my focus here is the structure of the underlying claims and how these might affect a comparative theory. cf. Feldman (1995a); Hurka (2001); and Kagan (ming).

⁷ On this point, I follow the lead of Shelly Kagan (Kagan, 2003).

cases, claims will need to be balanced against each other. How to do so is a matter of fairness. But developing that project must wait until Part III.⁸

Our project in Part II is more limited: What are the temporal constraints on claims? *When* does an individual have a claim? How does the basis of that claim affect *when* an individual has a claim? And how does an individual's temporal location affect the strength or existence of her claims?⁹

The answers depend, in part, on the type of claim to which those constraints apply. An overview of the different types of claims will therefore be useful. This is the subject of the rest of the present chapter. But first, a quick note on what claims are *not*.

5.1.1 WHY NOT RIGHTS?

When I mention claims, many have asked if they are rights; they are not. What is the difference? The exact difference is difficult to articulate, since at least the historical notion of rights “suffers from no small indeterminateness of sense,” and the confusion seems to continue today.¹⁰ This confusion — and the baggage that comes with it — is something I hope to avoid.

The main difference between claims and rights is that rights are usually understood to be trumping, in some way, of other values. When rights are violated, this is sometimes taken as ground for further action against the violator. By contrast, not all claims have this sort of strength. As we shall see, claims are one type of reason that gets pitted against others, and they may be denied. I therefore have tried to avoid any rights-talk so as not to mislead.

⁸ *infra* 174.

⁹ There are many other interesting questions about claims and time which I set aside. For instance, how does the passage of time affect the strength of a claim? And how does the time at which the claim is satisfied affect the value of that satisfaction?

¹⁰ *cf.* Griffin (2008). The quotation is from page 2. Griffin begins his book: “This book is prompted by the not uncommon belief that we do not yet have a clear enough idea of what human rights are” (1).

Most importantly, however, “rights” are not needed for a discussion of the sort I am pursuing. Claims and rights may turn out to be related; indeed, I suspect they are. But the relation may be a complicated one, and is beyond the scope of this thesis.

5.1.2 WHY NOT ENTITLEMENTS?

Others have asked if claims are entitlements; they are not. Rather, entitlements are like a peculiar sort of claim: ones that are rooted in an agreement, like a promise, contract, or law. Some entitlements matter morally, others may not. By “claim,” I intend the sort that do matter morally — that is, those whose denial or unfair satisfaction has disvalue — but determining which “entitlements” are of what sort is a matter for another time.

5.2 TYPES OF CLAIMS: AN OVERVIEW

There are two important distinctions to make between different types of claims. The first distinction is between primary claims and derivative claims. *Primary claims* are claims to objects directly. That is, primary claims do not depend on other claims. Examples of primary claims include moral desert claims and first-order entitlement claims. By contrast, *derivative claims* are claims held in virtue of their necessity for the satisfaction of other claims. For example, if an individual has a primary claim to health, then they might have derivative claims to those resources needed to stay in good health.¹¹ Or if one makes an agreement to buy some rice, then one has a primary claim to the rice, and a derivative claim to the company producing the rice. These examples are rough, but I hope they convey the distinction. Another way to describe the distinction is that the claim basis of a primary claim is *independent* of

¹¹ I provide this only as an example. As will be seen, I think that claims to health are actually derivative claims.

the other claims an agent has, whereas the claim basis of a derivative claim *depends* on the other claims an agent has.

The second distinction is between three broad types of claims, distinguished by what sort of reasons count as a basis. The three types are: Desert Claims, Claims of Convention, and Proximity Claims. Most of the claims which we speak about in ordinary language can be understood as one of these three. There may be additional types of claims, but I will ignore this possibility since our goal is not to provide a complete theory of claims.¹² There is already plenty of work to be done in understanding the basic three argued for here, and even our primary focus will be more limited to only desert claims.

Table 5.2 is a partially completed table listing the different types of claims and their properties. It is partially completed because, as noted, the goal here is *not* to offer a complete theory of claims, but to outline enough of the relevant features to enable our discussion about how some claim types work with respect to time.

The purpose of this chapter is to provide an overview of what these three types are, how they work, and to provide some comments to suggest that they are defensible (though I will not provide such a defense here). I will take each in turn.

5.3 DESERT CLAIMS

The first, and possibly most central, type of claim is a desert claim. Like other types of claims, it has three components: a subject, an object, and a basis — some reason in virtue of which the subject should get the object. The difference between desert claims and other types of claims is that the basis is restricted to some relevant fact about the *subject*, for example, his moral worth, or those things that might contribute to it, like his character or his actions.¹³

¹² As noted before, that is work for another time. To develop and defend such a theory here would bring us too far afield from the questions at hand.

¹³ According to Blackburn (1996, 102):

Claim Type	Agent	Object	Basis	Basis Type
DESERT (moral)			Properties of agent and/or relation to object <i>only</i>	
Moral Desert aka “worth claims”	Includes persons	Well-being	Worth of agent (goodness/badness, virtuousness/viciousness, etc)	Primary
Harvest Claims	Any productive agent	Object produced	Having caused the object to exist	Primary
Need Claims	Any agent with a worth claim	Objects	Needed to achieve well-being claimed	Derivative
CLAIMS OF CONVENTION			Depends in part on institutions and social norms	
Conventional Desert		Object / Prize / consideration	Properties of agent AND how they meet the criteria established by conventional norms, institutions, and practices	
Contractual Claims			Some sort of agreement, promise, etc.	
Property Claims				
PROXIMITY CLAIMS			Depends in part on agent’s proximity to object claimed	

Table 5.1: Types of Claims

The use and abuse of desert are ubiquitous in ordinary discourse, which can lead to confusion regarding the concept under discussion. For example, commonly heard desert ascriptions include: “Volunteers deserve to be rewarded.” “Wall Street deserves to be punished.”¹⁴ “Good papers deserve good grades.” “The fastest runner deserves to win the race.” But these so-called desert ascriptions are really making assertions about claims of two different types. The first two ascriptions, about volunteers and Wall Street, seem to track a notion of ‘karmic justice’. Underlying this notion of desert is the idea that things go better when an individual “gets what’s coming to them.” The second two ascriptions, about grades and races, seem to track a different notion of desert, one which is dependent not only on some feature of the individual (the quality of their work, their speed), but also on the conventions of the institutions involved, for the veracity of their claims.

For purposes of this thesis, I will limit the discussion to what is usually called *moral desert*. Moral desert is akin to the first sort of desert ascription made above, about volunteers and Wall Street.¹⁵ As noted, underlying this notion of desert is the idea that it is better (in a morally relevant sense) when individuals get their “just deserts”, or are subject to “karmic justice”. Put another way, moral desert

the term [‘desert’] denotes the idea that how a person ought to be treated depends on a fact about their actions, or their character: a proportioning of happiness to virtue, or unhappiness to vice. Treating persons as they deserve is then the exercise of justice. Problems include the distinction between moral and non-moral desert, and the question of whether the variations of treatment have a fundamental place in ethics, or whether they are to be justified, for instance on consequentialist grounds.

cf. Feldman (1997, Part III); Griffin (1986, Ch. 12); Kagan (1999); Kleinig (1971). For an interesting discussion of structural similarities between desert and virtue and vice within a consequentialist framework, see Hurka (2001).

¹⁴ “Or at least denied those bonuses which, we say, had always been more than they *deserved*.” They may never have deserved them — or at least that much — even though they may have been ‘entitled’ to them by their contracts. (Whether they had an entitlement claim to them in virtue of their contracts is a separate question.) For an explanation of the difference between desert and entitlement, see Griffin (1986, 255) and Kleinig (1971, 74)

¹⁵ Kleinig (1971) explains the concept of desert as distinct from the moral version. Griffin (1986, 254-257) offers a discussion of the “narrow” version of desert close to what I have in mind. Of course, moral desert might still come in many flavors, e.g. a victim may deserve remuneration for their injuries (“damages”). These may turn out to be derivatives of the principle — or something else entirely.

is roughly the principle that one should get something (good or bad) for a reason — and not just any reason, a reason directly tied to the goodness or badness of the agent as reflected through their character, actions, or other morally relevant feature.¹⁶

Claims akin to the second sort of ‘desert’ ascription made above, about grades and races, might also matter morally. I will discuss these briefly towards the end of the chapter. I set these apart because the claim bases are not strictly about the agent and/or the agent’s relation to the object, but also rely, in part, on the existence of the norms or agreements governing the institutions involved.¹⁷

Another confusion regarding desert is that we often hear desert claims made about things not traditionally considered as agents. For example, “The environment deserves protection.” Or “Great art deserves admiration and preservation.” Whether these are even appropriate subjects for desert claims is a separate question, one which would bring us too far afield. I will set these aside. Since my aim is an axiology about what is good *for people*, I will focus on the type of desert claims that are applicable to persons.

Finally, there might be thought to be two types of moral desert claims, distinguished by their object: claims to well-being, and claims to resources. I will discuss claims to well-being in the first section. I will then turn my attention to desert claims to resources.

5.3.1 MORAL WORTH CLAIMS: DESERT CLAIMS TO WELL-BEING

The primary type of moral desert claim is a claim to well-being in virtue of some morally relevant feature of the agent.¹⁸ That is, the primary type of desert claim (hereafter ‘moral desert’) is a claim where the object is well-being and the basis is

¹⁶ Note that the principle does not yet support specific measures for rewarding or punishing individuals.

¹⁷ *infra* 108.

¹⁸ *cf.* Kagan (1998, 54-59).

some morally relevant feature of the agent. Such morally relevant features might include, for example, the kind of thing the agent is (e.g. a person), or their character, intentions, thoughts, desires, decisions, or actions. The question of whether the satisfaction of this type of claim matters is, at heart, a question of whether it is better that agents be rewarded (made better-off) or punished (made worse-off) in accordance with their worth as indicated by these morally relevant features. The better they are with respect to these relevant features, whichever those relevant features turn out to be, the “more deserving” they are, and so, the thought goes, the better it is when they are better-off. Similarly, the worse that an individual is with respect to these relevant features — the “less deserving” they are — then the less good it is when they are made better-off.¹⁹ Some proponents of desert also believe in *retributive desert*. Retributive desert is the principle that, in some instances, when an individual is bad enough with respect to these morally relevant features, it is worse for them to be made better-off, and may be better for them to be made worse-off, or even for them to suffer. Retributive desert might be thought of as “punishment for the sake of punishment”, as contrasted with the use of punishment for instrumental purposes, such as protecting the public (in the case of incarceration) or as deterrence for other would-be evil-doers.

What are these “morally relevant features” in virtue of which an agent is deserving of well-being? Traditionally, the relevant criteria were often thought to be facts about the agent over which the agent has “control” or for which an agent may be said to be “responsible”.²⁰ These were generally taken to be some combination of an agent’s decisions, actions, intentions, or character, when such features might appropriately be ascribed to the agent (and not, say, to an evil scientist using the agent as a puppet). More illuminating, perhaps, is that certain features are irrele-

¹⁹ Kagan conceptualizes this as a “discount” on the value of their well-being as contributing towards the good. (Kagan, 1998, 55)

²⁰ For discussion of the difficulty of using these criteria, see e.g. Nagel (1979, Ch.3); also, Zimmerman (2002). For an argument against the traditional view, see Feldman (1995b, 175ff.).

vant to how well-off an agent deserves to be, like the color of one's hair, or skin.²¹ The reason normally given for excluding these is that they are facts about the agent over which an agent has no control. But this is not a particularly good distinction: it is possible to alter one's hair color (often very realistically), and even to alter one's skin color (albeit, to a limited extent).²² The better distinction is that features like hair color or skin color neither contribute to an agent's moral "worth", nor do they command a certain type of response.²³ By contrast, the aforementioned facts about an agent that are thought to matter — namely, actions, decisions, intentions and character — are thought to be indicative of an agent's moral worth.²⁴ At the least, they are thought to command a certain type of appropriate reaction, namely praise or blame.²⁵ But another way to describe this idea is that these are features that make an agent more or less worthy of whatever response the features command, from praise and blame to actual rewards and punishments.

Carving out the exact space of these morally relevant features would take us too far afield from our present task. Instead, I want to focus on a structural feature of this sort of desert, one that will be useful in our larger task of evaluating the Temporal Comparison Thesis. It is the feature that helps explain why, when faced with the need to reject Temporal Proximity Irrelevance or accept that there are no claims, we should choose the former. But first, I need to explain another structural feature of this sort of desert, the difference between *Absolute Desert* and *Comparative*

²¹ There have been times throughout history when these considerations *were* thought relevant to determining what an individual deserved. I will assume these positions to be wrong.

²² To wit, the culture of spray tans and tanning salons; although bleaching creams have been used in Asia for quite some time to achieve the opposite result.

²³ One might object that an individual's *natural* skin and hair color are things over which an agent has no control, and this is the reason for the moral irrelevance of such features. But contrast this with a feature of an agent's character, such as work ethic or patience, which is thought to be relevant. In the latter case, where easy fixes are available to the agent to improve this character trait, the agent is often taken to be blameworthy for not pursuing the easy fix, and praiseworthy for improving their character in this manner. That an individual has no control over the *natural* features with which they are born does not seem to absolve them of the responsibility to improve upon them where possible when such features *are* morally relevant.

²⁴ Subject to appropriate conditions, usually when these are under the agent's "control".

²⁵ e.g. Strawson (2003).

Desert.

ABSOLUTE VS. COMPARATIVE DESERT

An important feature of the structure of a theory of desert claims to well-being is that they have both an absolute and comparative component. As Kagan observes: “An adequate theory of (moral) desert will need to include not only noncomparative principles, but comparative ones as well.”²⁶

Absolute desert claims (hereafter Absolute Desert, or AD) are claims that an agent has to a certain amount of well-being in virtue of some fact about the agent.²⁷ That is, there is a specific amount (or range of amounts) of well-being that an agent absolutely deserves. Call the amount of well-being an agent deserves to have their *desert level*.²⁸

There are two important points about absolute desert claims. The first is that they constitute an essentially noncomparative measure: the amount an agent deserves does not depend on the amount that other agents deserve.²⁹ However, AD may have comparative implications when, for example, different agents deserve different

²⁶ See Kagan (2003, 93ff) for a defense; also, Hurka (2003).

But Arrhenius raises an important objection. He suggests that, if one rejects retributivist desert, then there is little left for comparative desert to do: it will not change the rankings given by his theory of absolute desert (Arrhenius, 2007, 29-38). Further development of my theory will need to address this problem. First, as Arrhenius notes (but doubts), there may be work left where “there is an unequal distribution of merit and too much or too little welfare to distribute” (37). Second, absolute desert under the present account differs from Arrhenius’s in that it is not exactly a “fit” account: it does not say how to distribute in cases where there is happily more welfare to go around than the individuals deserve noncomparatively even in cases where all have equal absolute desert. Third, there is reason to suspect that comparative principles may make a difference in different-number cases even if they do not change rankings for fixed populations. (I refer the reader to the relevant note on 250.)

²⁷ AD also goes by the name “noncomparative desert”. I use the two terms interchangeably.

²⁸ This is sometimes referred to as an agent’s desert (as in the common phrase you “should get your just deserts”), but I will use the term *desert level* so as to avoid confusion with the principle, also called desert.

²⁹ What about when an agent deserves a *share* of the relevant object? In the case at hand, of claims to well-being, this seems to make little sense. That an agent *deserves* to be well-off, or have a certain standard of well-being, should not be affected by other people. But having a (noncomparative) claim to a *share*, rather than a particular object or amount, might make sense in a different context of claims to resources. I will touch upon this briefly in the context of harvest claims (5.3.3).

amounts. The second is that an agent's desert level does not depend on whether it may be satisfied.³⁰ If it cannot be satisfied, this is worse from the perspective of desert, but does not affect the agent's desert level.³¹

There are several ways that AD assigns desert levels. The simplest is that AD assigns exact desert levels. That is, there is a specific amount that a given agent deserves. From the standpoint of desert, things are best when the agent receives exactly this amount — no more and no less. Call this *Exact AD*.

Other views are possible. *Minimum AD* is the view that there is a minimum level of well-being that a given agent deserves; that is, there is a limit to how badly off an agent deserves to be, but no limit to how well-off they could be. *Maximum AD* is the opposite view: there is a limit to how well-off an agent deserves to be, but no minimum.³² *Region AD* is the view that there is not a precise amount of well-being that a given agent deserves, but a range of levels that they deserve. In other words, there exists both an upper and lower limit to what they deserve, but it makes no difference (from the standpoint of desert) what level within the range they actually receive. *Exact AD* is a precise version of *Region AD*, and is the common interpretation of how desert levels are assigned.³³

Unfortunately, AD cannot always be satisfied. As a result, *comparative* principles are needed to explain when (and whether) one way of violating AD is preferable to another from the standpoint of desert.³⁴ *Comparative desert* is thus an account of how one's desert levels should be satisfied in relation to the satisfaction of other

³⁰ Or, in the case of varied desert objects, whether there is a desert source.

³¹ Indeed, this is a feature common to all noncomparative claims on this toy theory of claims.

³² This might be an extreme version of the Retributive View.

³³ Modal versions of these views are also possible. One might argue that there is a limit to how badly off or how well off an individual *could* deserve to be. This seems to be what underlies views that reject the use of torture for any purpose, regardless of what the individual who is to be tortured has done or plans to do. The views described in the text, in comparison, are about how badly off or how well off an individual *actually* deserves to be.

³⁴ This is not to suggest that there might be other reasons to prefer one way of violating absolute desert to another, e.g. maximizing well-being.

agents' desert levels. To put it another way, "comparative desert is concerned with how people compare in terms of how they are doing relative to what they (noncomparatively) deserve."³⁵

For example, one comparative principle I find plausible (though I shall not defend it here) is Kagan's *Y gap view*.³⁶ He explains the underlying idea as follows: "comparative desert is perfectly satisfied when (and only when) the offense against noncomparative desert is the *same* for all relevant individuals."³⁷

JUDGING EQUALS & DESERT BASELINES

This brings us to the second feature of desert that I wanted to focus on: there are *desert baselines* for the relevant kind of agent. A baseline is the minimum desert level that an agent deserves to be at — the minimum amount of well-being to which an agent may be said to have a *claim* just in virtue of being that sort of agent. It seems intuitive, at least for human agents, that there *is* a *minimum* baseline: there is a bottom limit to what any agent deserves to suffer.³⁸

What is it about an agent that makes them of the relevant kind for moral desert? That is, what is it about an agent that gives them a claim to a desert baseline, that makes them worthy of this response? My toy theory is that the relevant agents are those capable of appreciating their level of well-being.³⁹ There are two ways this

³⁵ Kagan (2003, 107).

³⁶ See Kagan (2003) for a defense of the Y-Gap View (107ff.) and an argument against the ratio view (100-107).

³⁷ Kagan (2003, 107).

³⁸ Certain extreme versions of the Retributivist View might not agree. That they would does not affect my point.

I have recently learned that Feldman has also proposed that "there is some modest level of happiness that people deserve merely in virtue of being people" (Feldman, 1997, 194). (Many thanks to Gustaf Arrhenius for calling this to my attention.) I am particularly sympathetic to Feldman's suggestion that this is what makes the Repugnant Conclusion so repugnant: "People should not be forced to live such drab lives. They deserve better" (Feldman, 1995c, 210). But I share Arrhenius's view that including desert within the axiology is unlikely to aid in the avoidance of the Repugnant Conclusion (Arrhenius, 2003, esp. 232-233). Even so, it can help in one small way: as I argue in Chapter 11, an account which takes comparative desert seriously would allow for a context-sensitive neutral level of well-being (*infra* 242).

³⁹ It might be objected: what of an agent who can only appreciate others' well-being, but not

might be cashed out, and I think the argument for desert baselines can be made from either.

The first is that the relevant agents are ‘reason-responsive’. One way of reading this is that the relevant agents are capable of judgment-sensitive attitudes in the Scanlonian sense. Judgment-sensitive attitudes are “evaluative attitudes that are responsive to the reasons an agent judges themselves to have (e.g. belief, fear, anger, admiration, respect).”⁴⁰ Another in this vein is the Kantian understanding of people as ends unto themselves. On either, the idea is that the appropriate agents have aims of their own — or ends — to which they respond as reason-giving. Because of this, harm done to them must be justified. Similarly, failing to aid individuals suffering in this manner must also be justified (i.e. there is a duty to aid such individuals). For such agents to be below the baseline — to suffer — is bad not only because suffering is bad, but because they have a claim to not suffer in such a way that reduces their well-being below the baseline. In the absence of those features noted above that might cause their desert level to deviate from the baseline, to cause them to suffer below the baseline is somehow also denying a claim.⁴¹

their own? I think this misses the point, but would depend on the favored account of well-being. To my mind, being able to appreciate others’ well-being is to appreciate one’s own well-being — it just ties one’s own well-being to that of others’.

⁴⁰ Scanlon (1998, 20-23). N.B. Although Scanlon rejects what he calls the “Desert Thesis” (274), my understanding of desert differs from the view Scanlon calls the “Desert Thesis”. This is because Scanlon seems to understand desert in the retributive manner. He writes:

It is sometimes said that feeling guilty for having done something necessarily involves the belief that one should be made to suffer in some way for having done it. Similar claims are made about related notions of blameworthiness, resentment, indignation, and so on. Let me call the moral idea underlying such claims — the idea that when a person has done something that is morally wrong it is morally better that he or she should suffer some loss in consequence — the Desert Thesis. (274)

I, by contrast, take desert to be the principle that individuals deserve to be well-off in accordance with their worth. Whether doing something that is blameworthy affects an individual’s worth, so as to make them worthy of punishment in accordance with the Desert Thesis, is a separate question. Indeed, my own inclination, like Scanlon’s, is to reject retributive desert.

⁴¹ In this way, it can be seen that my view does not entail retributive desert, but is compatible with it. For it is merely to say, in the absence of those actions or intentions which might cause an agent to be deserving of suffering, such an agent does *not* deserve to suffer in this manner — they have a claim to a certain level of well-being.

This seems tied with the second approach, which claims that the relevant agents are any agents with ends. This casts the net wider: it seems that animals could have their own ends in seeking pleasure or pain (even though they might not respond to reasons). I do not find it *implausible* that animals could deserve a certain level of well-being in virtue of being sentient agents. But the nature of their well-being might be different, and so they might not be assigned the right kind of desert by the principle at hand. I will not pursue the matter further other than to note that for this reason, I favor the first notion of the relevant agents.

Why think that these features of the agents ground claims to a desert baseline? The answer, I suspect, must presuppose that this type of agent has value. Valuing these features entails a certain sort of response: appreciating that they deserve a certain minimum level of well-being, in the absence of contravening desert factors.

Is this adequate? My intention here is not to defend the choice of the appropriate type of “agent”, or what makes them valuable. This is a problem faced by all moral theories, not just desert. Rather, it is to provide a rough explanation of the kind of agent that I think fits the bill and, roughly, why this entails claims of this sort. The type of agent I have in mind is a moral one; that is, whatever agent it is that we understand to be bound by the theory of the good of which desert might be one principle.⁴² If such an agent has a *desert baseline*, it puts them on the map, so to speak.⁴³ This is the importance of desert baselines, and we can describe these agents as having a desert level in virtue of the kind of agent they are. Call this baseline level of desert D_A .⁴⁴

The amount of well-being that an agent deserves might deviate from D_A . Indeed,

⁴² In this, I agree with Scanlon that the relevant agents — which include humans, though perhaps not exclusively — are those of a type capable of “the particular kind of judgments involved in moral reasoning.” Scanlon (1998, 179).

⁴³ I have recently learned that Kagan makes a similar point. Kagan (ming, 19-20).

⁴⁴ Perhaps, from the perspective of desert, an agent deserves *any* level above the baseline, though none is better than another (Minimum AD). I find this less plausible than Exact AD, but will not argue against it here.

the view that there is such deviation is the view commonly described as Desert. Objections against desert have been raised by denying that we have the relevant sort of control or freedom, whether it be freedom in a metaphysical sense, or, more strongly, from moral luck.⁴⁵ But these arguments do not succeed against desert. What they call into question is not *whether* people can be morally deserving, but, rather, whether people can be *more or less* morally deserving. That is, they call into question whether there is a reason in virtue of which people deserve a desert level different from D_A .⁴⁶ Differences in desert level might turn out to require control or freedom, but having a desert level does not. That treating people *differently* in virtue of actions, intentions or character might be undermined by these objections is not an objection to desert, but to assigning different desert levels in virtue of those actions, intentions, or character. Why? It is an axiom of desert that *equally deserving people deserve equally*.⁴⁷

I have thus far gestured at the sort of defense one might offer in favor of desert baselines, the basic claim that an agent has to a minimum level of well-being simply in virtue of being the kind of agent that one is (here, a person). I am afraid that I do not have a clear argument to support this intuition; I must assume it to be true. A real defense of this premise would take us far beyond the present aim of this thesis.

However, I do hope that I might address a worry about whether there could be such a thing as claims to *well-being*. After all, well-being is something of an abstract place-holder for that which makes a life go better or worse which allows us

⁴⁵ Moral luck suggests that what appears to affect the goodness or badness of an agent (in the relevant cases) is actually luck, not the morally relevant feature of the agent. Nagel (1979, Ch. 3). *cf.* Zimmerman (2002).

⁴⁶ Alternatively, they might be taken to show that we can be more or less deserving on the basis of factors outside of control — that is, on the basis of things subject to luck. Many do not want to accept this conclusion.

⁴⁷ Kagan (1999). I agree with Kagan that this is the principle underpinning attachments to another candidate principle for a pluralist theory of the good, Equality.

to bundle various goods onto one scale. Might it not then make sense, the objector might worry, to say that these claims are really to *resources*, to those goods which together form the bundle that is well-being? I will argue that such a view is misguided, and in so doing further motivate credence in this most basic type of desert claim, by exploring an important connection that this type of desert claim to well-being has with another type of claim: claims of need. Mostly, I hope it will help alleviate fears about the strangeness of granting that an individual might have a claim to *well-being*. Claims of need have traditionally been understood as claims to resources, and, in particular, to those resources needed for survival. I will argue that claims of need, conceived in a Wiggins-style elliptical sense, are really derivative desert claims to resources in virtue of this more basic desert claim to well-being.

5.3.2 DERIVATIVE DESERT CLAIMS TO RESOURCES: CLAIMS OF NEED

In addition to desert claims to well-being, there are desert claims to resources. A resource is an identifiable object (e.g. an apple) or opportunity (e.g. consideration for a job). There are two types of desert claims to resources. The first type is a derivative claim; the second type a primary claim. I will discuss the first type here, and the second type next in 5.3.3.

This first type of desert claim to resources is derivative on desert claims to well-being. Where certain resources are necessary to achieve the deserved level of well-being, a deserving agent can have a claim on those resources. I will call this the *Derivative Claim Principle*:

Derivative Claim Principle An agent has a derivative claim to a resource, X, in virtue of a primary claim *iff* it is necessary that if the agent's primary claim

is satisfied, then the agent receives X.⁴⁸

Prima facie, this principle is plausible: that an agent has a primary claim to A — that there is a *pro tanto* reason why they should get A — seems to entail that there is at least a *pro tanto* reason why, if B is a necessary means to A, they should get B.

A familiar type of claim in the literature, namely, claims of need, are of this sort within our toy theory of claims. Since this is the type of claim which will be our central focus in the following chapters, I will take some time here to explain how claims of need are really of this sort, and address two concerns regarding this way of developing the theory.

To say that an individual “needs X” is to imply that X is necessary for the achievement of some purpose. Indeed, as Wiggins observes, this is the standard reading:

Something that has been insisted upon in most analytical accounts of needing is that needing is by its nature needing for a purpose, and that statements of need which do not mention relevant purposes [...] are elliptical [...] for sentences that do mention them.⁴⁹

On the view that I’m presenting, the purpose of a need claim is clear: the achievement of that level of well-being to which an individual has a primary claim.⁵⁰ But there might be concern that this view takes an unnecessary further step: why not just have primary claims directly to the needed resources in virtue of the need?

⁴⁸ I have endeavored, in so far as possible, to follow the style used to express claims of need given in Wiggins, 1998 for reasons that will soon be clear.

⁴⁹Wiggins (1998, 7).

⁵⁰ I set aside the question of whether the locution of “need claim” applies only to those derivative desert claims which are in virtue of an agent’s desert baseline, D_A , or if it also applies to those derivative desert claims which are in virtue of an agent’s desert level more generally (which might deviate from D_A , if such deviation is possible).

This might have been the problem that Wiggins was suggesting (though my intention here is *not* a textual analysis of his theory) when he wrote that there was something about the elucidation of need claims as having a purpose that is correct, and something about it that is wrong. He argued that this type of analysis of claims of need — as being for some purpose that requires stating — captures the “instrumental sense” of the word need. But, he insists, there is an “absolute sense” of the word need, and that is the normative sense. While absolute need claims can be equivalently written as instrumental need claims whose purpose is the avoidance of harm, it is unnecessary to do so. The concept, he contends, already contains it, and “one does not have to supply again what is already there.”⁵¹

But a note by Anscombe, albeit taken from a slightly different context,⁵² might help to illuminate why it is important to take claims of need as derivative. She, like Wiggins, does not take claims of (absolute) need to be elliptical; but, unlike Wiggins, she does not appear to take them to be normative without making a further assumption:

To say that [an organism] needs that environment is not to say, e.g., that you want it to have that environment, but that it won’t flourish unless it has it. Certainly, it all depends whether you *want* it to flourish! as Hume would say. But what “all depends” on whether you want it to flourish is whether the fact that it needs that environment, or won’t flourish without it, has the slightest influence on your actions.⁵³

⁵¹ Wiggins (1998, 10).

⁵² Though, Wiggins drew my attention to it by quoting it in his discussion (Wiggins, 1998, 7 fn. 10).

⁵³ Anscombe (1958, 7). Anscombe’s discussion is about the transition from “is” to “needs” in the context of the impossibility of the transition from “is” to “ought”. She continues on to observe that, in the case of people, that something is or ought to be needed seems to have a normative connotation: “*that* such-and-such ‘ought’ to be or ‘is needed’ is supposed to have an influence on your actions” (7). So, this creates a weirdness that makes such claims dubious, unless there is some sort of connection between needs and wants. She observes:

Certainly in the case of what the plant needs, the thought of a need will only affect action if you want the plant to flourish. Here, then, there is no necessary connection

This is the further assumption I wish to make clear by making claims of need derivative on claims to well-being. That there is a reason to give an organism that which it needs depends on whether there is a reason to promote the flourishing of that organism. Similarly, if claims of need are to have the “special argumentative force” that Wiggins ascribes them, there must be a particularly strong reason why that organism needs to flourish in some particular way. For that an agent (or organism) needs a particular object in order to flourish does not yet get it a *claim* to that object — unless it is the case that that agent (or organism) already has a *claim* to flourish, or to the level of flourishing that such an object provides.

In other words, for need claims to have their “special argumentative force”, it is not simply enough that an organism flourishing is a good thing. That the organism’s flourishing is a good thing would provide a *pro tanto* reason for that organism to get that which it needs, since this is a necessary condition for its continuing to flourish. But this is a eudemonic reason to give the organism that which it needs, and so is quite different from the type of reason given by claims. Such eudemonic reasons also exist for giving organisms that which will make them better off. If claims of need rest merely on the value of an organism flourishing, then they lack

between what you can judge the plant “needs” and what you want. But there is some sort of necessary connection between what you think *you* need, and what you want.
(7)

Wiggins seems to miss this point. Though he opens his work with a lengthy discussion lamenting that need claims are often controversial despite both the pervasiveness of such claims and their importance, he concludes:

These people are representing, and the meanings of their words commit them to representing, that we simply can’t get on without more rods, more reactors, or more animal experiments or whatever. [...] the thing that is controversial in what is said is not the necessity for the avoidance of harm to human beings — it is precisely the fact that such avoidance is not a controversial purpose that lends needs claims their *prima facie* special practical and argumentative force —, but the claim that our prior adherence to this end *commits us* to have more roads, more fast reactors, more animal experiments or whatever. What is controversial is the claim that there is no alternative.
(Wiggins, 1998, 9)

Wiggins is right that the controversy often centers around the availability of the alternative. But it seems that there is also room for controversy over which harms are necessary to avoid, and the relative importance of avoiding such harms hold as compared to the cost of avoiding them.

special force when competing with these other eudemonic reasons. The special force of need claims comes from their status as *claims* — as something *owed* to a particular agent. And they must be derivative, since, for an organism to have a *claim* to that which it needs, it must first have a *claim* to the flourishing made possible by the object needed.

This point can be seen more clearly by considering the derivation of a claim of need in a case where claims are generally thought to play an important role, that of the need for an organ transplant.

1. Olivia has a (primary) claim to well-being D_A in virtue of being the sort of agent she is (i.e. in virtue of her worth). (*Assumed*)
2. Olivia's (own) organ is such that, whether through poor health or premature failure, the organ does not enable Olivia to achieve well-being D_A . (*Assumed*)
3. An agent has a derivative claim to a resource, X, in virtue of a primary claim *iff* it is necessary that if the agent's primary claim is satisfied, then the agent receives X. (*Derivative Claim Principle*)
4. There is an available organ, organ A, that would enable Olivia to achieve well-being D_A if she receives it as a transplant. (*Assumed*)
5. There are no other available organs. (*Assumed*)
6. It is necessary that if Olivia achieves D_A , she receives organ A. (*From 2, 4, and 5*)
7. Olivia has a (derivative) claim to organ A. (*From 1, 3, and 6*).

In other words, Olivia's claim to well-being, and the necessity of the donor organ to get that well-being, is what grounds her claim to the organ. Were she not to have a

claim to the well-being such an organ provided, it is unclear where her claim to the organ might come from.

There is a second concern, however, that might be raised, about why these claims of need are derivative on claims to well-being (or “flourishing”) rather than on claims to the avoidance of harm.⁵⁴ But what else is a harm if not one of the following:

1. allowing or causing someone’s well-being to fall below their claimed level of well-being
2. allowing or causing someone to be worse off than they otherwise would be
3. allowing or causing someone to be made worse off without justification

If the type of harm to be avoided is taken to be the first, then this is what I have in mind. It would not make a difference which way the formula is written, though using “harm” instead of the formulation in 1 is to introduce an unnecessary term that brings with it a certain amount of ambiguity. If the type of harm to be avoided is taken to be the second, it would not seem that this is strong enough to ground claims of need. This is because there are instances in which individuals are allowed to become worse off (e.g. by failing to make them better off) or even made worse off (e.g. by a policy, like taxation) that are not the sorts of harm that ground claims of need. If the type of harm to be avoided is taken to be the third, the allowing or causing someone to be made worse off without justification, there is still a question remaining: if the allowing or causing of harm is done by the failure or refusal to supply the needed resource, it seems similarly that the type of harm the person will befall must likewise justify the supplying of that resource for purposes of avoiding that harm. Otherwise, the justification for failing to supply, for example, a televi-

⁵⁴ Indeed, Wiggins’ formula is of this sort, where the purpose is “avoidance of harm” rather than “flourishing”, though “avoidance of harm can then be understood always by reference to a norm of flourishing relative to [the relevant circumstances]” (Wiggins, 1998, 13).

sion, might simply be that, while there is no reason *not* to provide one, there is also no reason *to* provide one other than the eudemonic reason that the individual will be worse off for it. This, again, does not constitute adequate grounds for a claim of need. To ground a claim of need, it must be the case that the harm — the being made worse off — is one of the sort against which they have claim. But to introduce this notion of “harm”, though perhaps rhetorically useful, obscures the theory.

5.3.3 DIRECT DESERT CLAIMS TO RESOURCES: HARVEST CLAIMS

A final type of desert claim to resources is distinguished by a particular kind of desert basis: what the subject has done to produce that resource. For example, if Adam took the time, care, and effort to grow a vegetable garden without the help of others, then Adam (and Adam alone) would have a desert claim to the harvested vegetables in virtue of the work he put in to grow them.⁵⁵ Since without Adam’s work, there would not be any vegetables to harvest, he deserves to have them.⁵⁶

What about if Adam didn’t create the resource, but added to its value? In some such cases, Adam will have generated a new resource by adding to the value of the original one; indeed, this is the manner in which he has added to its value. For example, suppose Adam had fertilized wild mushrooms growing in a field. The mushrooms are useful as food. Fertilizing these mushrooms causes them to grow both larger and more quickly. As a result, Adam has added to the value of these growing mushrooms, because they now produce more. This could be restated as: Adam has created an additional resource, namely, amounts of mushroom that would not have otherwise existed.

Other cases are more difficult. Suppose that Adam has developed a new method

⁵⁵ If trees were the sort of thing that could be deserving, one might think that the tree also has a desert claim to the harvest. I set this possibility aside.

⁵⁶ This claim could be overridden. Others might have claims of different sorts on his vegetables, in virtue of their need or through prior arrangement.

for extracting the nutrients from the mushrooms. This method does not extract any more nutrients from mushrooms than other methods; rather, it is more efficient. As a result, it takes less effort to get the nutrients from the mushrooms. What does Adam here have a claim to? One might think he has a claim to the time he has saved by using this method. But what claim does he have to the time others save by using this method? And can this extra time really be understood as a new “produced” resource? I must set these questions aside. The cases we will focus on in the next chapter are far simpler, and they present challenge enough.

But what if Adam discovered — rather than helped produce — the resource? The answer seems that it might be similar to the difficult cases of adding value: until the resource was discovered, its value was untapped. So Adam has a claim to some of that value. Delineating the exact extent of his claim is not needed for our purposes.

Finally, what if the work were done not only by Adam, but by a group of which Adam is a part? This is a type of desert claim where the noncomparative claim might be to a *share*, rather than a particular object or amount. Indeed, when we say Adam has a claim to the harvest he produced, we do not say that he deserves a particular amount of vegetables. Rather, we say that he has a claim to those particular vegetables that result from his labor, however many there happen to be.⁵⁷ Similarly, if the production of the harvest was a group effort, then it seems that each member of the group deserves a share of the particular harvest thereby produced.

5.4 CLAIMS OF CONVENTION

In addition to moral desert claims, there are also *claims of convention*. These are claims to specific goods or resources (e.g. house, land, car, or compensation in

⁵⁷ Whether it is best for him to keep them all is a separate question.

exchange for services rendered) that are in virtue of certain conventional norms. Claims to property are generally thought to be of this sort.⁵⁸

Conventional norms are norms that develop over time, usually for pragmatic reasons. Some philosophers believe that most of justice (with the exception of a few core principles) actually developed in this way. Some of those philosophers take the line that such norms of justice arose out of agreement of one sort or another that was reached for the benefit of those agreeing.⁵⁹ Others, like Hume, argue that they simply arose out of habit: people fell into certain patterns because they worked, and then the patterns became convention, but without explicit agreement or even discussion on anyone's part.⁶⁰

Whatever their origin, there remains an important question: if such claims are developed for pragmatic reasons, why must we include them in a theory of claims, the satisfaction of which is taken to have *intrinsic* value? To do so would appear to run the risk of double counting: these sorts of claims developed, in part, to maximize benefits accruing to individuals. That is, their value is *instrumental* in nature, not intrinsic. Thus, it may seem that we should not value the satisfaction of such claims *intrinsically* because their value is merely in their capacity to promote well-being, and so the value of their satisfaction (or dissatisfaction) will have already been counted in the distribution of well-being.

But it is also not clear that they should be excluded: for norms to be effective, they must be given value. That is, even though the claims that they generate are artificial (i.e. created), those claims, once generated, are very real.⁶¹ As a result, the

⁵⁸ e.g. Hume (1978, 3.2 SBN 484f).

⁵⁹ Actual contract theorists, like Hobbes, are of this sort. This is not to be confused with hypothetical contract accounts meant to *justify* moral principles.

⁶⁰ "Two men, who pull the oars of a boat, do it by an agreement or convention, tho' they have never given promises to each other. Nor is the rule concerning the stability of possession the less deriv'd from human conventions, that it arises gradually, and acquires force by a slow progression, and by our repeated experience of the inconveniences of transgressing it" (Hume, 1978, 3.2.9 SBN 490).

⁶¹ Their strength, however, may vary according to the legitimacy of the conventions. I will not pursue this matter here.

harm of denying a conventional claim is also very real, and is a harm in addition to the object (or well-being) so denied.

One might worry instead that this harm is merely a matter of expectations. In October 2010, the state of Arizona stopped funding certain types of organ transplants due to a budget deficit. Many patients who had been counting on state funding for soon-needed transplants were left in the cold. The legislature's decision has, sadly, resulted in the dissatisfaction of the moral desert claims — both primary and derivative — of several individuals who needed transplants urgently.⁶² However, at least some of the outcry seems to be not about the denial of those claims, but about the dashed expectations:

“To basically renege on what you promised was [going to] be a chance at life is a very, very bitter indictment of the ethics of the Legislature,” says Arthur Caplan, head of the Center for Bioethics at the University of Pennsylvania.

Caplan calls the reversal “awful” behavior because Arizona is going back on a covenant it made with its patients, and because these are patients for whom time is critical — patients who spent months, some years, thinking they were covered.

“They then stop trying to raise money, stop trying to see what Uncle Fred might be willing to give them,” Caplan says. “They don’t have the bake sale. They don’t make the appeal in church.”⁶³

It is obvious that there is something important about expectations: dashing them can be instrumentally bad. But the *instrumental* disvalue of dashed expectations cannot be the only bad that the critics are worried about. If I foolishly have the expectation that the guy at the mall dressed as Santa will stuff my Christmas

⁶² Unfortunately, as is often the case, the underlying claim landscape is complex.

⁶³ This is a real case (Robbins, 2010). Thanks to Olivia Bailey for drawing my attention to it.

stocking with the money I need to pay for the urgently needed transplant, and so make no alternate plans, that the mall Santa will dash my expectations is (for me) very instrumentally bad. Yet, this is *very* different from the Arizona case. In the Arizona case, the expectations were generated by a convention, unlike my foolish expectations which were generated by my fancy. As a result, the expectations in the Arizona case also have normative force: there is something more than instrumentally bad about them being dashed.

Still, the objector would be incorrect that it is expectations that do the work. It matters whether the convention generates a *claim*, not whether it generates expectations. The objector might think that the expectations in the Arizona case matter because (a) the system had been supporting transplants; (b) they were led to believe the system supported the transplants; (c) they were led to believe they qualified to receive them; and (d) such expectations were not only reasonable based on precedent, but perhaps also morally (given their need). Since the convention generated these reasonable expectations, these individuals seem to have a claim derived from those expectations, the denial of which is unfair.

By contrast, what about a racist old white man in the 1960s/70s who expected that his children would be able to attend school (supported by his taxes) in a black-free environment? His expectations seem to meet similar points as the transplant case, except for (d). Has he been treated unfairly when his children are forced to go to an integrated school, given his expectations?

I'm inclined to think not. He may have expected that his children would be able to attend school, supported by his taxes, in a black-free environment — and it is reasonable for him to have these expectations given the convention. But these expectations were generated by a bad convention. Bad conventions do not generate claims. Without a claim, he has not been treated unfairly — even though his convention-generated expectation has been dashed. By contrast, the Arizona

transplant patients had convention-generated claims — not just to a particular organ, but also to the particular funding that had been promised by the state. That their expectations have been dashed is instrumentally bad; that their claim has been denied is intrinsically so.

Does this mean that Arizona cannot change its policies? It seems it should be able to. This suggests that a further interesting argument for including claims of convention into the toy theory may come from cases where conventions transition. Claims of convention are needed to explain why the 2010 Arizona transplant patients had claims denied, when those in 2012 do not have claims of this sort. But this leaves much undecided. For example, what claims does someone who falls just at the boundary have — who discovers, just prior to the policy change, that they need an organ?

Claims of convention are particularly interesting because they encapsulate a temporal asymmetry. Claims of convention cannot exist before the convention does (although they may bootstrap each other into existence). But if the argument above is correct, once they exist they *do* have value. Moreover, there comes a point after the convention has ceased when there are no longer claims generated by that convention, and perhaps also when claims already generated by that convention cease to have value. The question then, arises, whether anyone can have a claim under a given convention, or if some are necessarily excluded by their location in time.⁶⁴ Unfortunately, I must leave these questions for another work. While I may occasionally refer to claims of convention in what follows, it will be mostly to set them aside.

⁶⁴ A similar question could be asked about conventions with respect to space.

5.5 PROXIMITY CLAIMS

The final type of claim are proximity claims. A proximity claim includes in its basis the claimant's proximity to the object claimed. The inclusion of this type of claim is likely to be controversial, and so its defence is a large undertaking which this thesis will not attempt to complete. It will require first the rejection of a principle called Proximity Irrelevance — the principle that an individual's *mere proximity* to a resource is irrelevant to the existence or strength of claim upon that resource. This is the project I turn to next, with respect to the relevance of temporal proximity.

AGAINST TEMPORAL PROXIMITY IRRELEVANCE

We turn now to the problem of temporal distance as it relates to claims to resources. There are two parts to this problem, but this thesis will address one: how claims to resources might vary when the relata (i.e. subjects, objects, or bases) are located at different times, whether owing to having merely partial temporal overlap or no temporal overlap.¹ The second problem, which I will not discuss, is about how the time of the satisfaction of a claim might influence the value of its satisfaction.

This chapter aims to attack a common assumption about the relevance of temporal proximity to both the determination of claims to resources and to those claims' comparative strengths. The assumption, called *Temporal Proximity Irrelevance*, is this:

Temporal Proximity Irrelevance An individual's temporal proximity to a resource is irrelevant to the strength of her claim on that resource.²

¹ How might a base exist at a different time than the subject or object? I'm not entirely sure. I leave it open only as a possibility.

² In practice, it must also assume *Spatial Proximity Irrelevance*, the spatial analog of Temporal

We begin this attack by considering a *very* simply growth model in which there is only one sort of resource, or capital, namely, a magic mushroom.³

The Growing Mushroom There is only one sort of resource, a magic mushroom.

This mushroom grows over time. Once it is depleted, it will not grow back.

There are no other resources in this world.⁴

The world, as in many of our examples, contains a fixed population of k generations, each generation with n individuals. That is, there are k generations, each with the same number of people.

We may now ask two questions:

1. How should we divide up the mushroom between the generations?
2. Who has a claim to what?

The first question, on the toy theory of claims that we have been developing, will depend, in part, on the answer to the second question.⁵ This chapter focuses on the second question.

I first consider the answer that claims are to the initial mushroom. Second, I discuss what some might consider an obvious solution: that claims are to a share of the new mushroom produced each year (“mushroom excess”). Third, I consider the solution that claims are to shares of the mushroom aggregate (that grows over time) considered atemporally. I argue that none of these three options are consistent with

Proximity Irrelevance. For now, assume that the individuals in the world are located in a relevantly close region such that the spatial distance between contemporaries relative to the mushroom is negligible.

³ Thanks to John Broome for drawing my attention to this thought experiment.

⁴ I take resource to mean a material commodity of some sort. Resources need not have intrinsic value, though they may, in unusual circumstances, have intrinsic value. Rather, resources are usually desired for their instrumental value — that is, their capacity to be put to use in increasing those things which are of intrinsic value.

⁵ Note that this first question seeks an all-things-considered judgment which the toy theory of claims does not yet offer. Nor will offer, at least not until it is combined with a theory about the disvalue of complaints as compared to the value of well-being. This is part of the work of Part III (*Infra* 174).

Temporal Proximity Irrelevance. Finally, I consider and reject the solution that no one has a claim to the mushroom. As a result, I argue that we ought reject Temporal Proximity Irrelevance.

6.1 CLAIMS ARE TO THE ORIGINAL MUSHROOM

The first potential answer is that everyone has a claim to the *original* mushroom. So conceived, the claim is to the mushroom *as it exists* at time $t = 0$.

What might support such claims? There are two potential claim bases in virtue of which the claims to the original mushroom are equal.

First, individuals' claims might be equal in virtue of their equal need: the people in this world need the mushroom to sustain themselves (i.e. to attain the level of well-being to which they have a claim). As discussed in Chapter 5, these need claims themselves derive from primary claims to well-being. So, this answer makes three assumptions: each individual has an equal claim to well-being, the manner in which a mushroom contributes to an individual's well-being is the same for each individual in this population, and either that the amount of mushroom assigned in this manner does not exceed the amount needed or that, where the mushroom exceeds the amount needed, being equally comparatively deserving of well-being grants equal comparative derivative claims to the mushroom.

Second, they might have equal claims in virtue of the fact that they are all equally deserving of the mushroom, or, rather, equal claims in virtue of the fact that they are all equally not deserving of the mushroom: no one has done any more than any of the others to bring about the mere existence of this mushroom. So, they have equal claims (or, at least equal comparative claims).⁶

⁶ On the difference between noncomparative claims and comparative claims, see 5.3.1. This hints at another possible answer: that *no one* has a claim to the mushroom in virtue of desert. As I suggested earlier (5.3, page 89), combining need and desert claims under one category might help resolve this difficulty: If humans *deserve* a certain minimum level of well-being (in virtue of the kind of thing they are), then they will have a claim to the amount of resource needed to satisfy

Other types of claims, like conventional claims such as property, do not as yet figure into determining the claims. There is nothing about the mushroom that establishes it as one person's property (except, possibly, proximity of some sort, which would conflict with TPI), and there are as yet no conventions in place. One might object that the mushroom is somehow communal property. But there does not seem to be anything in the existence of the mushroom to establish this: if there is nothing about the mushroom that establishes it as one individual's property, it does not follow that it is communal property.⁷

But whichever of these bases establishes the claims, each person in the history, regardless of their temporal proximity to the mushroom, has the same claim. Why? This is because this method of claim assignment assumes *Temporal Proximity Irrelevance*. That is, it presupposes that temporal proximity to a resource is morally irrelevant: temporal proximity does not matter in determining claims to that resource. On this view, the strength of a claim is not affected by the temporal location of the claimant relative to the resource so claimed. As a result, temporally later claimants have equally strong claims as temporally earlier claimants to the mushroom at $t = 0$. If mere temporal proximity *did* matter, then there would be differences in either the claims or the strength of the claims that these individuals belonging to different generations would have to the mushroom.

But this leads to a problem. In our model, the mushroom will grow. (This is taken as given.⁸) As a result, the slice of mushroom to which individuals at later

that (and no more). This last suggestion may turn out to map well a sort of sufficientarianism.

⁷ For comparison: undiscovered ideas are generally not considered to be communal property.

⁸ This model was developed to be similar to an exogenous growth model, although it differs from the neoclassical model in that it does not require the input of labor. An exogenous growth model takes technological progress as coming from outside the system. That is, the decisions made about consumption and savings do not affect technological progress (and the resulting increased productivity). As a result, decisions about how to divide the mushroom do not affect the manner in which the mushroom grows. In future work, I hope to discuss the more complicated case where decisions about how to divide the mushroom do affect how the mushroom grows. This is akin to endogenous growth models, which take technological progress as part of the system. Barro et al.

times have a claim will have time to grow before they are able to enjoy it, and so members of later generations, it seems, will enjoy more mushroom than members of earlier generations. How much additional mushroom they have to enjoy depends on the growth rate of the mushroom. Denote the particular mushroom piece claimed by person $p_{i,j}$ (the j th person in the i th generation) as $m_{i,j}$, and the size of that slice of mushroom at time t as $m_{i,j}(t)$ (for times prior to consumption, $0 \leq t \leq t_{\text{consumption}}$). Then, assuming that the mushroom has an instantaneous growth rate of r , the size, at time t , of the piece of mushroom that individual $p_{i,j}$ has a claim to is:

$$m_{i,j}(t) = \left(\frac{M_0}{k \cdot n} \right) e^{rt} \quad \text{for } 0 \leq t \leq t_{\text{consumption}} \quad (6.1)$$

where M_0 is the size of the original mushroom at time $t = 0$.⁹

The difference between the size of the mushroom slice at $t = 0$ to which the later individual has a claim and the size of that mushroom slice at the time of the later individual's existence, t , is $\Delta m_{i,j}(t)$ (i.e. $m_{i,j}(t) - m_{i,j}(0)$). Call this the *mushroom excess at time t* . But do later generations have a *claim* to the full amount that has grown from their share? That is, do they also have an exclusive claim to the mushroom excess, the additional amount of mushroom that their slice has at the time of their consumption?

The answer depends, in part, on how the resource is understood. Is the mushroom excess a new resource? Or part of the original one? We'll consider each in turn.

6.1.1 MUSHROOM EXCESS: A NEW RESOURCE

If the mushroom excess is considered to be a new resource that has appeared in addition to what remains of the original mushroom, then whether the members of

(2004, 61f).

⁹ For greater detail, please see Appendix A.1 on page 263.

later generations have an exclusive claim to the mushroom excess of their slice of original mushroom depends on whether time has brought new claims. If time has not, then the claims to this new resource are determined in the same manner as the claims to the original resource were. Those claims to the original resource were grounded by the Temporal Proximity Irrelevance Principle. Since this principle is not asymmetric with respect to time, the claims of all generations to this new resource will be equal.

Of course, time might have brought new claims. There are two claim types which have such a potential to be introduced: desert claims, and conventional claims. I'll discuss each in turn.¹⁰

The first type of new claim which may have been introduced is a desert harvest claim.¹¹ This type of claim might plausibly be thought to arise within the context of an individual life: if individuals are temporally extended, as we have assumed,¹² it would allow individuals to reap the benefit of their decisions regarding what to do with their share of the mushroom (consume less earlier or more later). Individuals who choose to consume at a later time would have additional mushroom which had grown out of their original slice – mushroom excess to which they have a claim in virtue of having produced it by delaying their consumption. Any resulting inequalities would then be a result of individual choice.¹³

But the analogy from desert claims within a life to the intergenerational context does not work: the later generation cannot be said to have done anything to deserve the additional mushroom, since they did not have the option of consuming it earlier. The only individuals who might be said to have acquired a desert claim

¹⁰ There is a further type of claim that might be thought to exist, namely ones stemming from risk. If it is unknown whether the mushroom will grow or shrink, then it may affect how claims to the initial resource – and any subsequent growth – are divided. I must set this complication aside.

¹¹ Recall that harvest claims are direct desert claims to an object that the agent has produced in virtue of the agent having caused the object's existence. (5.3.3, *supra* 107)

¹² Not everyone holds this view. How and whether desert claims exist, given those views, is a separate question. I set it aside. *cf.* Parfit (1984).

¹³ As before, I here leave aside the question of whether this type of desert could be justified.

to the new portion of mushroom that exists at time t are those individuals from earlier generations who respected the claim of later generations. We will address this complication in a moment. For now, it is clear that desert harvest claims have not been introduced for the generation at time t (or subsequent generations) that were not present at the initial time of decision, $t = 0$.

The second type of new claim which may have been introduced since time $t = 0$ is a conventional claim. But which convention we should follow is precisely the question at issue: until a convention is established, and, perhaps more strongly, determined to be just, we cannot use convention as a reason.¹⁴

Thus, it seems clear that time will not have brought new claims with it, at least not for those members of later generations who exist at the same time (or later times) as the mushroom excess. Ignoring the complication of additional claims for members of generations past, the claims to this new resource of excess mushroom must be determined in the same manner as the claims to the original mushroom. That is, everyone, across all time, has an equal claim to the excess piece of mushroom.

These claims to the mushroom excess at time t present two problems. First, the pieces of mushroom excess to which members of later generations have a claim will, if left, grow mushroom excesses of their own. Second, members of past generations, who live before time t , will not be able to access that piece of mushroom excess to which they have a claim.¹⁵

The first problem, that the slices of the excess saved for future generations will

¹⁴ Whether the mere establishment of a convention is enough to provide reasons to respect it, apart from details about how it was established or whether it is just, is a question for another time.

¹⁵ John Broome writes: “Also, there is no fixed amount of mushroom to be distributed among the generations. The less the earlier ones get, the more there is to be distributed. I’m not sure equal division makes good sense.” (Personal correspondence.) In a way this is right: it would seem that anyone’s claim to a piece of mushroom excess depends on whether the excess exists. I will address this momentarily. In a way, however, it is not. In assigning claims, we can ask what would be plausible given different histories of growth.

grow and create new excesses of mushroom, is just a reiteration of the original problem. It will be solved in a similar manner, for the same reasons: under Temporal Proximity Irrelevance, each person, regardless of their location in time, has an equal claim to the excess thereby grown.

The second problem creates a different difficulty. Members of earlier generations cannot directly obtain or make use of that to which they have a claim on this view. This is because what they have a claim to (on this view) is the time-locatable resource, the slice of mushroom excess at time t , and time-locatable resources at times later than an individual's existence are inaccessible.

What is to be done, then, with the slices of mushroom excess claimed by but inaccessible to members of earlier generations? One possibility is to reassign claims to that piece of mushroom excess to the members of the present and future generations. This is a reiteration of the original problem, but with the mushroom excess at time t , rather than with the original mushroom at time $t = 0$.

A second alternative is to allow for early collection: members of earlier generations would collect their claims early by trading their pieces at a later time t for additional mushroom at the earlier time at which they exist. But there is an obvious difficulty with this method, namely that it would affect the mushroom growth: an early harvest of a piece means that that portion will not have the opportunity to grow, thereby reducing the total mushroom excess by the amount of excess that would have come from the prematurely-collected mushroom piece. As a result, it might be said that earlier generations who engage in the early collection practice effectively destroy the time-locatable resource — the mushroom excess at some later time t — to which they would have had a claim. Thus, early collection does not allow them to collect that to which they have a claim; rather, it destroys it. Worse, it might also be said that members of earlier generations who engage in early collection also destroy the time-locatable resource to which others had a claim. As a result,

the earlier generations who engage in early collection might be thought to acquire a negative desert claim, according to which they no longer deserve the amount of later excess mushroom because they prevented its existence. More strongly, they might now deserve *less* than the others because they have effectively destroyed some of the stock. Early collection seems an impossible method: on either of these readings, it destroys the claim that justified the early collection in the first place by inhibiting mushroom growth.

But does this then entail that it is impossible for those members of earlier generations to obtain that to which they have a claim? On this view, with claims to the time-locatable resource assigned in this manner, then the answer is yes: there is no way to obtain it without preventing the existence of what would have been claimed (and thereby preventing the existence of the claim itself).¹⁶ But it is also not a good thing, on this view, that members of earlier generations cannot obtain that to which they have a claim, for it is a bad thing for claims to not be satisfied.

Is this a problem for this method of assigning claims? It seems that it might be. That an individual's claim cannot be satisfied does not entail that they do not have one. In other cases where individuals' absolute claims cannot be satisfied, we must appeal to comparative claims. Comparative fairness (as contrasted with absolute fairness) might take the burden of impinging on the claims of members of earlier generations for the satisfaction of the claims of members of later generations to also constitute an element of unfairness.

To prohibit earlier people from collecting early is to allow the claims of later people to impinge on the satisfaction of earlier people's claims while denying the earlier people the same privilege. This casts doubt on whether it really is the case that members of earlier generations ought not to collect early.

¹⁶ This, however, does *not* justify the asymmetry in our valuation of progressive inequalities discussed earlier, where later generations are better off than earlier generations (above 3). This is because that an inequality cannot be fixed does not entail that it would not be better if it could be fixed.

Is this a problem? It seems strange that this problem would be built into our very way of talking about claims over time. It is not a superficially contingent fact that these temporally earlier individuals cannot obtain that to which they have a claim; it is a deep, structural one. Only a change in the way in which time works could resolve their claim.

A further problem might shed some light. Something about the manner in which this discussion of claims has continued seems peculiar: there seems to be a nonexistence problem, similar to a nonidentity problem, but for claims.¹⁷ In particular, the existence of anyone's claim to the mushroom excess depends on the earlier generations having allowed it to grow. This can be seen by considering a case in which the mushroom does not grow: there are no further resources beyond the original mushroom to which individuals might have a claim. This would be a rather simple case, and the solution under discussion — that claims are to a share of the original mushroom at $t = 0$ — might seem a strong, unproblematic contender.¹⁸

This fact, that the existence of claims to the excess depends on the existence of the excess, might lead one to conclude that there is an easy resolution to this problem: as long as the first generation¹⁹ saves in such a manner that later generations receive a slice that has *the same size* as their share of the original mushroom, their claims will have been satisfied. One might reason to this conclusion as follows:

¹⁷ Recall that we have already set aside, for the present discussion, the problem of non-identity of the persons with the claims. It will turn out not to affect our conclusions, and is a useful simplifying assumption.

¹⁸ What if the mushroom were shrinking? The problem, then, might be seen more clearly: if claims are assigned to equal shares of the original mushroom, then the members of later generations — just in virtue of their temporal proximity to the resource — will have a claim to a resource of lesser value. I will discuss a similar problem with respect to the growing mushroom in 6.1.2. While the case of the shrinking mushroom would be more straightforward, in a sense, there would not be the question of what counts as a “new resource”. Since the problem of new resources is one that arises even in a simple case, I have decided to use a growing mushroom instead of a shrinking mushroom.

¹⁹ A similar, though slightly more complicated, argument could be run using ‘earlier generations’ in place of ‘the first generation’.

First, each generation has a claim to their original mushroom slice, $m_{i,j}$. Claims to subsequent mushroom excess, $\Delta m_{i,j}(t)$ depends on the mushroom having so grown. If the mushroom does not grow in this manner, there will be no additional resource beyond the original mushroom to which anyone may have a claim. If the members of the first generation consume (collectively) more than M_0/k (i.e. $1/k^{th}$ of the original mushroom), this cannot be said to violate later generations' claims to any mushroom excess that might have grown. Hence, the only claim of later generations that this action of the first generation violates is their claim to their share of the original slice of mushroom. But, if the first generation saves in such a way as to allow later generations to receive a mushroom slice which is at least the same size as the slice of the original mushroom to which they had a claim, it would seem that later generations have received an appropriate substitute for that to which they had a claim.

This seems a happy solution. But while it provides a suggestion for an alternative method of assigning claims (discussed shortly²⁰), it does not solve the problem at hand. Later generations can object that this trade is unfair. That to which they have a claim, the slice of the original mushroom, was more valuable than what they receive through the trade: the slice of the original mushroom had growth potential. This is so even when treating the mushroom excess as a new resource, since the slice of original mushroom's potential was precisely in generating this new resource. This objection is also consistent with the problem of the nonexistence of claims raised a moment ago, namely, that later generations do not have a claim to the excess that would have grown (but does not). This is because the potentiality of the thing to which they had a claim (the original mushroom slice) was part of the original resource.

But then, to divide the claims this way is to distribute a more valuable good

²⁰In 6.3, beginning on page 128.

— an investment with returns — to some individuals and not others, just in virtue of where they exist in time. This is certainly not to make an individual's temporal location irrelevant to the determination of what they have a claim to.

It also raises a question about what claims to resources attach to: the particular temporally-locatable resource, or to the value which that resource might generate. In asking whether individuals' claims are to *particular* commodities or to the value of those commodities, we effectively ask this question: does it matter that those with equal claims to a resource receive equal portions of that resource, or portions of the resource which are of equal value?

This echoes a concern raised above: this claim-structure necessarily biases against those who live earlier. That it does so is due to a deep structural fact about our world and the nature of time, however contingent the nature of time might be. Is such a claim structure fair? More importantly, is this consistent with TPI?

That time — and *mere* time — creates a bias on this version of claim assignment is a serious problem for the Temporal Proximity Irrelevance principle which grounds it. Those who would take the principle seriously must consider another option.

6.1.2 MUSHROOM EXCESS: STILL THE ORIGINAL RESOURCE, JUST BIGGER

I mentioned above²¹ that, in addition to the question about claims, there was also a question about how to conceptualize the mushroom excess. I considered first the understanding that the mushroom excess was a new resource, different from the original mushroom. But there is an alternative: the mushroom excess is part of the original resource; the original resource has just grown.

Conceiving of the mushroom excess in this manner causes the assignment of claims under TPI to immediately run into the problem just described: it is to

²¹ *Supra* 118.

assign, from the outset, claims to resources that have different value in virtue of an individual's location in time relative to the time needed for that resource to manifest its potential. The claims that were established on this view were in virtue of the individuals having equal claims to the well-being such a resource might provide, that is, to the resource's value. For time to be directly relevant to the value of the resource to which an individual has a claim, therefore, flies directly in the face of TPI. Again, advocates of TPI must reject this method of claim assignment.

6.2 CLAIMS ARE TO THE MUSHROOM EXCESS GROWN IN A GIVEN YEAR

A second option which many find appealing is to say that each individual has a claim to a share of the mushroom excess of their time. That is, the original mushroom is held in common, and, each year, the mushroom excess — the accrued interest, if you will — is divided between those who live during that year.²² Proponents suggest that this method is appealing because it grants everyone the same amount of mushroom given our assumptions about a constant rate of growth and constant population-at-a-time. As a result, they suggest, it is temporally neutral, since it treats everyone equally.

But for proponents of TPI, this option is something of a nonstarter: it makes the time at which one exists directly relevant to the claim one has on the mushroom. As a result, it is *not* consistent with TPI.

This can be shown by considering the two assumptions upon which its equal distribution rests: the rate of growth, and the population-at-a-time. Vary either of

²² The choice of one year as the period for distribution is arbitrary, and done only for sake of simplicity. Might it affect the claims? If we maintain our simplifying assumption that individuals have either complete temporal overlap or none at all, both with each other, and with a given year, then it will not. But it will affect claims if individuals vary with respect to when they are born and when they die, and if there is partial overlap —with either other individuals, or with periods so measured.

these two features, and the resulting distribution will not be equal. More importantly, however, the resulting *claims* will also not be equal. Holding population-at-a-time constant, those who live when the growth rate is slower will have a claim to less mushroom, since there will be less mushroom excess in those years. Holding growth rate constant, those who live when the population-at-a-time is greater, will have a claim to less mushroom, since there will be more people among whom it must be divided. And vice-versa. This result is not merely because of changes in growth rate or population, but because claims are linked to the growth rate or population *of the time when* one exists. Consequently, on this view, claims are directly related to the time at which one exists. This should have been obvious, since the position might have been restated thus: each individual has a claim to a share of the mushroom excess produced in a year in which they live. This may be a plausible solution, but it is not one available to those who defend Temporal Proximity Irrelevance. Indeed, it flies directly in the face of it.

This option is also inconsistent with Temporal Proximity Irrelevance in another way, given our assumption about the number of generations being finite, since there will be the original mushroom remaining after the final disbursement to the final generation. Who has a claim to this?

One answer might be no one. But, as will be discussed in 6.3.2, this is not particularly plausible, especially on an egalitarian view.

Another answer is that it should be divided evenly between the final generation. But if this is so, it seems like the members of generation g_{k-1} have a claim to consuming part of the original mushroom endowment, since this may be done so as not to reduce the amount that those in g_k receive. To prohibit this early consumption is to have built-in to our account of claims a bias against earlier generations. While less of what g_k receives will be from the mushroom excess to which they have a

claim, there will still be more than the original amount to which they had a claim available through consumption of the remainder of the original mushroom. A similar argument could be made for the members of generation g_{k-2} , g_{k-3} , and so on.

This is a similar problem to the one encountered by the first option of assigning claims to the original mushroom. Another solution must be found if TPI is to be maintained.

6.3 CLAIMS ARE TO A SHARE OF THE HISTORIC MUSHROOM AGGREGATE

One other such solution is this: claims are to a *share* of the historic mushroom aggregate. The historic mushroom aggregate is:

Historic Mushroom Aggregate the total amount of mushroom to actually exist in a given history, considered atemporally; i.e. the sum of all mushroom throughout time in a given history.

There is a counting problem here about how to consider pieces of mushroom that “exist” at multiple times. I will assume that mushroom pieces are temporally extended for sake of simplicity. This should not affect the results in this section, though I will not pursue that matter here. Since our interest is in the aggregate of *consumable* mushroom (whether or not it is actually consumed), one way to “count” the mushroom aggregate is then at the point of consumption. For example, if all consumption is delayed until the last time, T , then the size of the mushroom at any particular time t before T will be given by:

$$M(t) = M_0 e^{rt} \tag{6.2}$$

where M_0 is the size of the original mushroom and r is the mushroom's instantaneous rate of growth. The historic aggregate — counted as the total amount consumed — will be the amount of mushroom available for consumption at time T , $M_0 e^{rT}$.

Of course, consumption will happen throughout history. The timing and amount of consumption affects mushroom growth, and, in turn, the amount of mushroom available to be consumed (i.e. the size of the mushroom aggregate). So the historic mushroom aggregate is not fixed. But this does not entail that claims cannot be to equal shares. Rather, it means that everyone must consume in such a manner that everyone receives an equal share.

Recall that the Historic Population is the total number of people who ever live (for a given history). Equal claims thus entitle each individual to an equal share of the historic mushroom aggregate, or $\frac{\text{Historic Mushroom Aggregate}}{\text{Historic Population}}$. On this view, an individual does not have a claim to a *particular* slice of mushroom. Rather, their claim is to a slice of mushroom of the appropriate size.

As before, this claim is in virtue of one, or both, of two claim bases: each individual has a claim to the mushroom in virtue of their equal need (since they need it to sustain themselves), or in virtue of being as deserving as any other individual. That is, they have an absolute claim to the mushroom in virtue of their need, and these claims are equal in virtue of their needs being equal. They also have a comparative desert claim as strong as anyone else's to the mushroom. Given Temporal Proximity Irrelevance, each individual will have the same claim regardless of their temporal location.

For a given mushroom of original size M_0 , which has an instantaneous growth rate of r , the amount that each individual would have a claim to is:²³

$$m_{i,j} = \frac{\int_0^T C(t) dt}{k \cdot n} \quad (6.3)$$

²³ I here, and throughout, assume continuous time.

where $m_{i,j}$ denotes the amount of mushroom to which $p_{i,j}$ has a claim, T is the last time at which the mushroom exists, $\int_0^T C(t) dt$ gives the size of the historic aggregate (the total “consumed”), and k and n are the number of generations and the population of the generations respectively.

This appears to be an appealing solution: on the face of it, it would appear to satisfy Temporal Proximity Irrelevance, and it assigns claims in a plausible manner according to need, or desert more generally. It also captures the intuition that where individuals are equally deserving, they deserve an equal *share* of the available resource (rather than a particular amount).

But this view also encounters trouble. I divide the discussion of the trouble into two parts: first, an objection from interference to growth from natural causes; and, second, an objection from interference to growth from a malicious agent. The responses to these problems are in tension with each other, with the result that this view must be rejected.²⁴

Interference from Natural Causes: Mushroom Blight

Consider the following case:

Mushroom Blight (Unexpected) Suppose that there is a world with a magic mushroom which grows on a common field with an instantaneous rate of growth r . There are no other resources in this world. On the view under discussion, an individual has a claim to consistently consume an amount of mushroom, m , at each time during their life. Individuals in this world consume that to which they have a claim.²⁵

Then, suddenly, at some time partway into the history of this world ($t_{\text{blight}} > t_0$), there is an unexpected mushroom blight which inhibits the mushroom’s

²⁴ Many thanks to Olivia Bailey for her help in getting the structure of this argument to flow.

²⁵ Each individual would then have a lifetime consumption of mushroom $\frac{mT}{k}$. The size of m given these conditions is provided in Appendix A.1.4, page 266.

growth by destroying a piece of the mushroom of size m_{blight} . In doing so, the blight reduces the size of the historic mushroom aggregate from what it otherwise would have been. As a result, there is no longer enough for each individual to consume m at each time during their life.²⁶

How do we treat claims on this case? According to the view that claims are to equal shares of the historic aggregate, claims have shifted. Since the historic mushroom aggregate is now smaller, the size of mushroom slice to which any individual in this history has a claim is also smaller. But since the blight occurred at some time into the history ($t_{\text{blight}} > t_0$), and since individuals consume the amount to which they have a claim, those who existed before the unexpected mushroom blight will have consumed what they originally had a claim to. Given the size of the new historic mushroom aggregate, we must now say that these members of the earlier generations have consumed more than that to which they had a claim.²⁷ That is, at the time of consumption before the blight ($t < t_{\text{blight}}$), the individuals consumed the slices to which they had a claim. Later, at times after the blight ($t > t_{\text{blight}}$), it turns out that they consumed more than that to which they had a claim. This seems a strange result.

The strangeness of the result depends, to some extent, on the underlying metaphysical views. On a simplistic view that combines eternalism with a closed future,²⁸ there is only one actual future which as always existed, and so it would have been true at times before the blight ($t < t_{\text{blight}}$) that there would be a blight at t_{blight} . As a result, there is, truly, only one historic mushroom aggregate upon which to base claims, and it was always the case that those individuals who consumed mushroom

²⁶ This is shown in A.1.4, page 269. The amount of the decrease will depend on both the size of m_{blight} , as well as on the time at which the blight occurs. The earlier the time of the blight, the greater the reduction in the historic mushroom aggregate.

²⁷ The size of this discrepancy will depend on the size of the reduction in the historic mushroom aggregate, which depends on both the size of the reduction in the growth rate, as well as on the time at which the blight occurs.

²⁸ See 7.1.2.

before t_{blight} were consuming more than their fair share (assuming that they consumed the full amount of mushroom that they believed they had a claim to without the blight). Thus, the problem would not appear to be a particularly large one for this method of claim assignment. It may have *seemed* that the earlier individuals were consuming that to which they had a claim at times before the blight, and that then, at t_{blight} , the size of their claims *changed*, since there was no way they could have known that the blight would occur. But this is only an illusion which can be explained away once epistemic considerations are taken into account. Proponents of this method of claim assignment can then, with relative ease, “bite the bullet” and say that, indeed, members of pre-blight generations consumed more than their fair share.²⁹

Indeterminism presents a different challenge. But a full discussion of the problems posed by indeterminism would bring us too far afield at this point, and so I set it aside for now. The most interesting problem it raises is how to deal with *genuinely chancy objects*. In some ways, this may not be too dissimilar from dealing with *genuinely vague objects*, should such things exist. Both create serious problems for proponents of TPI.

The case of (Unexpected) Mushroom Blight thus does not pose a serious challenge to the view that claims are to shares of the historic mushroom aggregate, provided that its proponents are willing to bite the (small) bullet and allow that those who live before the blight have consumed more than their fair shares. Indeed, they must do so to remain consistent with TPI; otherwise, location in time relative to t_{blight} will have a direct impact on claims.

²⁹ The question of whether these members are blameworthy for having done so is a separate question, but one which is readily resolved on views where merely having done the wrong action is insufficient to ascribe blameworthiness.

6.3.1 INTERFERENCE FROM MALICIOUS AGENT

Suppose that mushroom growth has been inhibited not by blight, but by an individual:

Mushroom Vandal Suppose that there is a world with a magic mushroom which grows on a common field with an instantaneous rate of growth r . There are no other resources in this world. On the view under discussion, an individual has a claim to consistently consume an amount of mushroom, m , at each time during their life. Individuals in this world consume that to which they have a claim.

In the middle of the night, at time t_{vandal} , a mushroom vandal creeps out onto the field, cuts off a chunk of mushroom, and destroys it.³⁰ This inhibits the mushroom growth in much the same way as in the case of the (Unexpected) Mushroom Blight.

Does this mushroom vandal thereby merely reduce that to which the others have a claim? Or does he destroy it?

The most plausible answer is that he has destroyed it. He not only made the others worse off — by reducing the amount they might actually receive — but he also denied their claims by destroying that to which they had a claim.

But this response is not available to those who would establish claims as a share of the historic mushroom aggregate. This is because the claims of other individuals were not to a particular piece of the mushroom, but to a share of the historic mushroom aggregate. The offender, in destroying a chunk of the mushroom, has reduced this historic mushroom aggregate. The amount by which he has done so will even be larger than the piece of mushroom he actually destroyed, given the potential of that destroyed piece to grow. But since it is the size of the actual historic mushroom

³⁰ Whether he destroys it by consuming it, or by burning it (with no benefit to himself), does not matter.

aggregate that determines the size of the piece to which individuals have claims, the amounts to which they have a claim is now smaller. For most individuals³¹, the mushroom vandal has not destroyed anything to which they had a claim; he has merely prevented its existence. In such a way, he has made them worse off. But he cannot be said to have violated their claims.

This is a particularly poignant problem in cases where the vandal is not a member of the first generation, but lives (and acts) relatively early on in the world's history. Suppose that the vandal lives during the third generation. His actions then affect, to a great extent, the amount of mushroom ever available.³² This, in turn, affects the size of every other individual's claim. But this yields two bizarre results: First, he, as before, cannot be said to have violated the claims of future generations.³³ Second, and perhaps, more strangely, he has now made it such that earlier generations, who consumed the amount that would have been their "fair share", have consumed more than that to which they had a claim.

This is an unacceptable result. Moreover, it is a problem common to theories that assign claims solely as claims to a *share* of an *aggregate*, where individuals may influence the size of that aggregate.

How might the proponents of this view respond? They might say that individuals have a claim to an equal share of the *potential* historic mushroom aggregate rather than the *actual* historic mushroom aggregate. Here, 'potential' must be understood as 'without agent interference'. In such a way, the mushroom vandal may be said to have violated others' claims: he prevented that mushroom excess to which those other individuals had claims from existing. This response also avoids the problem

³¹ I use this qualifier because the mushroom vandal must have destroyed a piece larger than the piece to which he himself had a claim, and, in so doing, did destroy at least a small sum to which others may have had a claim.

³² To see the amount affected, replace m_{blight} with m_{vandal} and t_{blight} with t_{vandal} in A.1.4, page 269.

³³ Note that this denial that he has violated the claims of future generations does not originate from the usual source of such a denial, namely, the non-identity problem.

of saying that the earlier individuals, who lived before t_{vandal} , consumed more than their fair share when they consumed what they had a claim to without the vandal's interference.

But this response is in tension with the response to the first case, of mushroom blight, considered above. In that case, the proponent of the view under consideration wanted to maintain that individuals' claims were determined by the size of the *actual* historic mushroom aggregate. Failing to say this makes the time at which an individual exists — before or after t_{blight} — directly relevant to the size of mushroom share to which that individual has a claim. This would be in violation of TPI, and so the proponent of claims to the shares of the aggregate needed to bite the bullet and accept that the claims were determined by the *actual* historic mushroom aggregate.

What is the difference between the two cases of Mushroom Blight and Mushroom Vandal? Can the proponents of the shares to the aggregate solution find a principled way to divide the line between when claims are determined with respect to the *actual* historic mushroom aggregate versus with respect to the *potential* historic mushroom aggregate?

Consider a case in which the inhibition of the mushroom growth is caused by an agent, but accidentally so:

Mushroom-Destroying Klutz Suppose that there is a world with a magic mushroom which grows on a common field with an instantaneous rate of growth r . There are no other resources in this world. On the view under discussion, an individual has a claim to consistently consume an amount of mushroom, m , at each time during their life. Individuals in this world consume that to which they have a claim.

At time t_{klutz} , a klutzy individual walks out onto the field, and accidentally crushes a chunk of the mushroom, thereby destroying it. This inhibits the mushroom growth in much the same way as in the case of the (Unexpected)

Mushroom Blight.

My own intuitions on this case are a bit confused. It seems as though the Mushroom-Destroying Klutz has violated others' claims, however unintentionally. But, at the same time, this inhibits mushroom growth in a manner similar to the mushroom blight: without intentional action on the part of any agent, the mushroom's growth has become inhibited unexpectedly at some point in history. There does not seem to be a principled way to draw the line between this case and the mushroom blight one. The only solution is ad hoc.

Therefore, the view that claims are to shares of the historic mushroom aggregate must be rejected.³⁴

³⁴ There is a further problem encountered on this view which I have chosen to exclude from the main body of discussion. It is this:

Recall that mushroom growth was assumed to be continuous, and that people were assumed to be temporally extended. The result is that during any individual's life, the mushroom is growing. But, at what time do we distribute the shares?

If shares are distributed at the start of each generation, then individuals might consume their mushrooms differently, with some consuming early and others waiting. Since the mushroom grows continuously, this entails the size of an individual's piece of distributed mushroom will vary depending on when they consume it. More importantly, it will also affect the size of the historic aggregate of mushroom in virtue of affecting a portion of that aggregate. This, in turn, affects the size of each individual's claim. This is because claims, on the view under discussion, are to a *share* of the *historic aggregate* – a share of the total amount of the mushroom considered atemporally – and not to any individual slice of that mushroom. Individual consumption patterns affect the size of the aggregate, and, as a result, the claims to a share of that aggregate. Those who refrain from consuming their piece of mushroom increase the size of the aggregate. But if not everyone consumes homogenously, then those who allow their mushroom to grow end up with a share of the mushroom aggregate which is larger than that to which they have a claim of the sort under discussion. This is a bizarre result.

This seems counterintuitive: by simply exercising self-restraint (or by being forced to delay consumption) an individual will generate more mushroom than that to which they have a claim because the mushroom in their possession will be larger than an equal share of the total aggregate. (Indeed, if everyone so acts, then the share will, at most, be the equal share.) Surely, the individual, whether through choice or coercion, who delayed consumption, should enjoy the fruits of that delay. But they will not have any additional claim to the good unless their waiting generated such a claim.

This observation suggests a potential resolution: it might be said that there is another type of claim which now plays a role, namely a production-type desert claim. This might work in one of two ways. First, they might *be* more deserving, in virtue of their restraint, and so deserve (absolutely) more well-being. As a result, they will also (plausibly) deserve more comparatively. That is, they will deserve a larger share than the others, since they will require more mushroom to satisfy their primary desert claim, and since they will have a larger derivative comparative claim. Second, they might have a harvest claim to the excess mushroom which they were able to generate by waiting. This seems a plausible way to resolve the difficulty, provided that it is not a problem that they have also augmented the original claims of others. A theory would need to

6.3.2 NO CLAIMS

The final option consistent with TPI is this: no one has a claim — any claim — to the mushroom.

At this point, I must make a confession: I am partly sympathetic to this view, which is not a particularly comfortable place to be in for someone trying to salvage something in the vicinity of teleological egalitarianism. There would not appear to be anything in the *mushroom* itself that would generate the existence of claims to it. It seems obvious enough that the mere existence of something useful does not generate a claim to it, as we saw when I introduced claims in 5.1.³⁵

Worse, I am afraid I do not have particularly strong or good arguments against this view. What is one to do?

The best I can hope is to try to establish whether this view, that there are no claims, is consistent with the egalitarian view under discussion. If it is, then the egalitarian has a choice: abandon TPI, or abandon claims and find their own way out of the arguments of Part I. If it is not consistent, then the egalitarian must abandon TPI. For those who have long since given up on telic egalitarianism, this constitutes part of an evaluation of my argument in Part I for its rejection.

One thing that is clear is that egalitarians must take the distribution of mushroom to matter. This is because, according to the telic egalitarian, it is bad for some

be worked out to determine how the new claim balances against those of the original sort. The claims-to-shares-of-the-aggregate only conflicts with this position: that *only* the individual who allowed his mushroom piece to grow has a claim to the extra mushroom. But denying this is not such a difficult bullet to bite; presumably a theory which balances the claims of that individual with the claims of others (stemming from their claim to a share of the aggregate) might match such a judgment.

There is a larger problem here, however, and it is this: the actions of those who consume early reduce the size of the claims of those who do not consume early. Even if those who delay consumption have a *stronger* claim to that which they produced, that they do not have a claim to the full amount that they have thereby produced is not within their control. Worse, it is not merely that it is not within their control — this seems less problematic in the case of natural causes, like natural disasters or mushroom blight — but it is in the control of another agent much like themselves.

³⁵ *Supra* 84.

to be worse off than others through no fault of their own. Since the mushroom is the only resource available in the world, and consumption of the mushroom makes individuals better off (and in the same manner), then if some individuals get more mushroom than others, they will be better off than those others. Since, according to the egalitarian, it matters how well off people are with respect to each other, the distribution of the mushroom matters. This is something that the egalitarian cannot avoid, nor should they wish to.

However, that the distribution of the mushroom matters doesn't yet tell us anything about whether there are claims. Nor has the egalitarian explained yet *why* it is a bad thing for some to be worse off than others through no fault of their own. But of course, they need not. One option is for the egalitarian to take the Fundamental Egalitarian Insight as basic. They could reason as follows:

1. Each individual has the same utility function (i.e. the mushroom contributes to each individual's well-being in the same manner). (*Assumption*)
2. There are constant returns to well-being of the mushroom. (*Assumption*)
3. The amount of well-being an individual has is directly proportional to the amount of mushroom that individual has. (*From 1 and 2*)
4. It is bad for an individual to be worse than others through no fault of their own. (*Fundamental Egalitarian Insight*)
5. Each individual should get a piece of mushroom the same size. (*From 3 and 4*)

This seems to work, and is a way for the egalitarian to avoid claims. But there are two problems with it.

First, it yields a view similar to that just discussed and rejected. According to Step 5, each individual should get a piece of mushroom the same size. This sounds

much like the view presented in 6.3, that claims are to a share of the historic mushroom aggregate (i.e. divide the actual total of the mushroom equally). This makes claims derivative on the badness, however. It is not the case that the inequality has disvalue because claims are or are not satisfied equally. Rather, it is the case that because inequality is bad, there are reasons why individuals should get a certain amount. This subjects the egalitarian not only to the arguments of Part I, but also to the objection in 6.3 if the egalitarian wants to take the further step of prescribing how much an individual should get.

The second problem is that taking the insight as basic runs into the problem of Part I, and does not seem to leave any way out. The refusal to explain *why* it is a bad thing for some to be worse off than others through no fault of their own seems to leave no room to explain whether and why some comparisons might matter more than others. Indeed, this is partly what generated the problem in Part I.

It seems that the egalitarian might do better to allow for claims, to take claims as basic, and to adopt the principle that it is bad for some to be worse off than others with respect to things claimed. This would combine the view that what an egalitarian really cares about is that those who are equally deserving (i.e. who have equal claims) get equally. When they do not, this is unfair — and constitutes the badness of inequality.³⁶

To take this route out, however, is to reject TPI: no other way assigning claims is plausible or consistent with the principle. This leads to the new problem of discounting in morality: how does time matter to the strength of claims? I turn to this next.

³⁶ In this sense, this view takes its inspiration from arguments made in Kagan, 1999, which argues that what egalitarians really care about is that equally deserving people get equally (combined with the view that individuals are equally deserving), and Broome, 1989, which argues that the badness of inequality comes from individuals being treated unequally with respect to their claims.

TEMPORAL PROXIMITY RELEVANCE: THE NEW PROBLEM OF DISCOUNTING IN MORALITY

In the previous chapter, I argued against the principle of Temporal Proximity Irrelevance (TPI), the principle that an individual's temporal location relative to a resource is irrelevant to the strength or existence of her claim on that resource. I argued, in particular, that this was the case for derivative claims to resources stemming from desert (including claims of need).¹

The question remains, however, in what way *is* an individual's temporal proximity to a resource relevant? To answer this question, I consider a simple case where only one individual has a (non-competing) claim to a resource. After addressing some preliminary assumptions, I turn first to a discussion of how the time and existence of an object can affect a derivative claim to that object. I then consider

¹ I did not address TPI with respect to claims on well-being. Nor did I discuss whether temporal proximity is relevant to conventional claims of any sort (including property claims, and conventional desert).

whether and how the time and existence of the claimant affects their derivative claims.

The goal of this chapter is not to completely develop a theory of how claims work relative to time, but to evaluate a few plausible constraints that might apply to derivative desert claims to resources. Since this is the only type of derivative claim I will be discussing in this chapter, I hereafter use the simpler phrase “derivative claim” to refer to derivative desert claims.

7.1 PRELIMINARIES

When do individuals have claims to resources? And how does this relate to the temporal location of a claim’s three relata — the agent, the object, and the basis? To answer these questions, I must first address a few preliminaries concerning assumptions I make about how temporally-located resources contribute to well-being, about time, and about the appropriate way to time index claims.

7.1.1 WELL-BEING & TIME

First, I set aside the problem of *when* an individual has well-being, and how temporally-located goods contribute to that well-being. Specifically, this excludes the problem of whether events that happen during times when an individual does not exist affect her well-being without making any particular time in her life go better. I do so by simplifying to *lifetime* well-being. I do not anticipate that this will substantially affect our results at this early stage in developing a theory of claims. In all the examples used, it will be assumed that the contribution made by a particular resource to an individual’s well-being are to their lifetime well-being.²

² This assumption should not come as a surprise; I have been making it since the beginning of this work (see Section 2.1).

7.1.2 METAPHYSICAL ASSUMPTIONS

The second preliminary is to stipulate the metaphysical picture of time which is being used for purposes of this discussion. When discussing problems of how normative principles apply over time, it is not surprising that a number of metaphysical problems arise. *Prima facie*, it would seem that certain ways of treating value or obligations depend on the metaphysical picture assumed.³ Indeed, the question of whether and, more importantly, how differing metaphysical views about time (and related problems) should affect our normative views is a fascinating one. Unfortunately, it is a problem I need to sideline for now.

My present approach is to begin with the simplest picture, one that combines eternalism with a closed future. By *eternalism*, I mean the thesis that all times are equally *real* in the sense that there are no times that have a somehow metaphysically privileged status (e.g. as presentists take the present to be). By a *closed future*, I mean that the future exists as it will exist. (This is something that might be thought to be entailed by an eternalist picture, but it is worth stating explicitly.) I do so to set aside the difficulty that facts are not necessarily eternal, since, on this view, they are eternal: if at any time *p* is true, then *p* is true at any time.⁴ My hope is that understanding how claims might work on this relatively simple view

³ The debate over free will and moral responsibility is just one instance where metaphysics are thought to be relevant to the normative picture.

⁴ This is not to take a view on the relationship *between* events that is implied by determinism or indeterminism.

That debate is not about time; rather, it is about how the events within time relate to each other. *Determinism* is, roughly, the thesis that, given some set of initial starting conditions and some set of natural laws, every event within the history is completely determined. Those natural laws might be merely large patterns of regularity (in the Lewisian/Humean sense), or something stronger. By contrast, *indeterminism* is, roughly, the thesis that, given some initial starting conditions and some set of natural laws, not everything in the history is determined (i.e. the history is underdetermined).

You could combine indeterminism with eternalism. What that entails is that the relationship between at least some events, A, and some events, B, is a matter of chance. That B followed A, rather than C, was not certain. To make use of the possible world heuristic, there is a very close world to the actual world in which C follows A rather than B, and there is nothing else between the two worlds prior to the time at which either B or C instantiates that is different.

This distinction is not an important one for our purposes. What matters is that, given our actual world, one of B or C follow A (but whether it was certain that it would given the starting conditions and the laws does not matter).

of the metaphysics of time might help illuminate how best to alter the theory to accommodate an open future and/or a limitedly open future.

One might worry, however, that the assumption of a closed future could generate problems on the theory of claims which I have developed, given that the primary focus has been on desert claims. This is because, historically, the question of desert has been associated with the problem of free will and moral responsibility, which in turn were thought to depend on the truth of determinism.⁵ And, while I have not assumed determinism, it is not uncontroversial that a closed future is compatible with indeterminism. So, it seems that this might pose a problem. If we assume, however, as I have, that each person in a history is equally deserving, the problem of the compatibility of moral responsibility with determinism will not affect our results. The problem of moral responsibility and its compatibility with determinism, as relates to desert, is about whether individuals may *differ* with respect to what they deserve, and not with whether they deserve a particular amount just in virtue of being the kind of thing they are (as I have assumed).⁶ For this reason, I set aside the problem of the compatibility of moral responsibility with determinism, though it remains an interesting and important question for further development of my theory.

7.1.3 TIME-INDEXING CLAIMS

Finally, I must include time-indices within the structure of a claim. That is, I must reformulate “A has a claim to X in virtue of Y” as:

A has a claim at time t to X at time t' in virtue of Y at time t'' .

I reformulate it in this manner because, though the three time indices may be related, it is not yet clear that they are the same.⁷

⁵ Watson (2003) offers a good collection of essays on this and related points.

⁶ 5.3.1, *supra* 97.

⁷ This formula is not dissimilar to the one offered by Wiggins (Wiggins, 1998, 7):

It is necessary (relative to time t and relative to the t circumstances c) that if

There is something more to be said, however, about the indexing of claims, and it concerns the first time variable t to which the claim itself is indexed. Is this necessary? On the view of time under discussion, one with a closed future (and a closed past), it might not be. If an individual has a claim at some time t to some object at time t' in virtue of some basis relative to time t'' , and all times are real and exist (including those times t , t' , and t''), then it will also be (tenselessly) true that that individual has a claim *simpliciter*.⁸ In this way, claims are like reasons: they are timeless, on this view.

It now should be clear why an indeterminate or open future might complicate the discussion: if the future is open, then it is difficult to see how it might be tenselessly true that an individual has a claim at some time in the future t . For, if the future is uncertain, then it might not yet be true (or false) whether that individual will have a claim at some future time t . For it may not yet be true whether that individual will need that object at that time. Nagel has a helpful suggestion on this point:

(... at t'') then (___ at t').

Whatever is expressed by “...” is the thing of concern for which “___” may be said to be needed. He notes a similar reason for the use of three time indices:

The three time variables t , t' , t'' , are all needed in the full form of the ‘need’ sentence as given without any ellipse or abbreviation (even though, evidently, they do not vary entirely independently). t'' need not be t' , because the goal mentioned in the antecedent may be achieved later than the prerequisite mentioned in the consequent. t need not be t'' , because the end-state, a child’s having good second teeth for instance, may create an earlier necessity that certain measures be taken. And t need not be t' either, because it may be necessary *now* (t) that certain measures be taken next year (t') for the subsequent (t'') good of the second teeth—even though it was not necessary last year (before t) that these measures be taken at t' . Before t there was a chance of dispensing with these t' measures, e.g. by taking proper or better care of the first teeth. (Wiggins, 1998, 7-8, fn12)

He continues: “It is well beyond my present scope to trace the relationship of the time indicators in this sort of elucidation to the tenses that appear in the various natural language forms that it seeks to elucidate.” This tracing, one might say, is my present purpose.

⁸ This might be thought to match our use of “Chris has a claim to using the boat today.” By this statement, one might mean one of two statements: “Chris has a claim_{today} to using the boat” or “Chris has a claim to using the boat_{today}”. What we mean by the statement is that Chris has a claim to today’s use of the boat, rather than Chris has a claim today to use of the boat generally. This seems to me, at least, to be the more natural reading. The objects of the tensed claim statement are what bear the time index, not the claim itself.

There are two ways of dealing with this difficulty. (a) One can say that in such a case there *is* a tenseless truth about the future and a derivative reason for action in the present, but a prospective agent has no way of knowing this except through the kind of evidence which would justify him in believing a future tense statement about the [need for an organ]; consequently he has reason to arrange [to accommodate this need] but cannot know it — or (b) one can say that the agent has no reason to promote an end until it is true that the reason-predicate holds or will hold of that end. Alternative (b) amounts to an interpretation of tenseless truth in any of the three major tenses. It is not an abandonment of the timelessness of reasons because although the conditions which confer a reason on an event have been conceded not to be timeless, the value which that reason embodies is still timeless, and can transmit its influence from future to past, once it is clear that the reason will be present.⁹

To put this into the language of claims, what is uncertain is whether the conditions for an individual having a claim will obtain. But if those conditions do obtain, then we may say that they have a claim *simpliciter* to the (time-indexed) resource in a way that features into the tenseless axiology. That is, the value of a claim's satisfaction transmits to the present time in the same manner as other value.

Given this, for purposes of this chapter, I will refer to individual's derivative claims *simpliciter* to time-indexed resources. Whether these formulations do in fact need to be revised for use under an open future (or open past) is a question that must wait. As we shall see, the simple picture is difficult enough for now.

⁹ Nagel (1970, 54-55).

7.2 TIME AND EXISTENCE OF THE OBJECT NEEDED

I now turn to the problem of exploring the relationship between the other two time variables in the claim formulation, and what conditions they place on a theory of claims. I begin by considering how the basis of a claim affects *when* there is a resource to which an agent might have a claim (and so, one of the conditions under which an individual may have a claim). To do so, I consider a familiar case where claims are thought to apply and where those claims are thought to constitute at least one part of fairness, namely the case of those in need of organ transplants.¹⁰

In that vein, consider the following case: Suppose that Olivia has a desert baseline of D_A .¹¹ That is, Olivia has a claim to a level of lifetime well-being equal to D_A . Suppose further that for Olivia to achieve this level of lifetime well-being, she must live until time t_{D_A} . A premature death — a death prior to t_{D_A} — would cause her lifetime well-being to fall below her desert baseline D_A .¹²

Unfortunately, Olivia suffers from a congenital heart disorder. There is a time, $t_{\text{organ failure}}$, which is before t_{D_A} , at which Olivia's heart will fail. If Olivia does not receive a new heart transplant by $t_{\text{organ failure}}$, then she will die. Her death will be premature, and will therefore cause her lifetime well-being to fall below her desert baseline, D_A . Hence, if Olivia is to have the level of well-being to which she has a

¹⁰ cf. Broome (1984a,b, 1989). Broome takes claims as given, setting aside the question of what grounds them (Broome, 1984a, 44). Our use differs in two ways: First, I take claims to have both noncomparative and comparative uses, whereas Broome's use is only in the comparative. Second, we differ as to what fairness requires with respect to claims.

¹¹ As discussed above, in Section 5.3.1, a desert baseline is the minimum desert level that an agent deserves to be at — the minimum amount of well-being to which an agent may be said to have *claim* just in virtue of being that sort of agent.

¹² This assumption is designed to make the following discussion simpler by removing alternative ways in which Olivia might achieve this level of well-being, say, by having a life that is shorter but sufficiently better at each time to achieve this lifetime level.

In practice, the reason that a premature death might prevent achievement of D_A could vary. In many cases, this will be due to the amount of joy she would have experienced were she to live a longer life. In others, it may be as a result of the inability to complete certain key life projects, vital to her well-being. As a result, the age at which suffering an early death would impact well-being with respect to the desert baseline might vary from person to person, and could be off-set (or at least mitigated) by other factors as noted.

claim, she must receive a new organ by time $t_{\text{organ failure}}$. Since she needs an organ in order to achieve the well-being to which she has a primary claim, it would seem that Olivia *needs* an organ in a claim-generating sense. But is the fact that she needs an organ sufficient to establish a claim to *an organ, any organ*? Or are there further conditions that must be met for her to have this sort of derivative claim to a resource?

I will begin answering this question by considering the nature of the resource which Olivia needs, namely, a donor organ to replace her failing one. To avoid complication, assume that a donor organ is an organ which is available for transplant that is not currently being used by anyone.¹³

Should such an organ exist, it would seem that Olivia would have a claim to it. There might be others who also have a claim to the donor organ, in which case there would be competing claims. Who should ultimately get the organ would depend both on what fairness demands when balancing competing claims, and on a theory about how the value of fairness and the value of well-being trade off. But for now, I set this problem of competing claims aside.

One condition on the donor organ is clear: the mere existence of a donor organ is not enough for Olivia to have a claim to it. The donor organ that is required must exist at the right *time*. Consider, for example, a history in which an organ becomes available, but only after $t_{\text{organ failure}}$. Does Olivia have a claim to *this* organ? The answer is no: she does not need it. Obtaining this organ will not make it possible for her claim to D_A to be satisfied; this time-locatable resource is of no use precisely because it does not exist at the relevant time. Since the basis of a need claim is that the resource will provide or enable the level of well-being to which the individual has a primary claim, need claims cannot attach to objects which are not useful in

¹³ It is a separate question whether an individual like Olivia could have a need claim to an organ that is currently in use as an organ (as opposed to in use as a decorative conversation piece) by another person. By stipulating that “donor organ” is one which is not in use, I hope to exclude this possibility—for now—to avoid unnecessary complications.

this manner.

Indeed, any donor organ to which Olivia might have a claim must exist, at minimum, at time $t_{\text{organ failure}}$. It may exist before that (indeed, it would probably be less nerve-racking for Olivia if it did). She will also need it to exist for some time after $t_{\text{organ failure}}$, until time t_{DA} , if it is to help her achieve a lifetime well-being of D_A .¹⁴ It would not seem to particularly matter to Olivia whether and for how long the organ exists outside the bounds of the time when she needs the organ, provided that this does not affect the quality of the organ. So, for simplicity, assume that the organ quality is unaffected by the duration of time it exists, and the duration of time in which it is in use.¹⁵ In this case, the time of the object to which Olivia would have a claim, t' , seems to be the range of times from $t_{\text{organ failure}}$ to t_{DA} inclusive. This is the object that Olivia needs, and might have a claim to, if it exists.

There is an interesting feature of this object to which Olivia might have a claim. It is temporally *extended*. There is an alternative way in which to describe this picture, which might be revealing. Instead of saying that there is *one* object to which Olivia has a claim, namely an organ that exists for that specific range of times, there are *many* objects to which Olivia has a claim, namely each of the temporal parts (or stages) of that organ during that range of time. A temporal part is defined as follows:

Temporal Part “ x is a *temporal part* of y iff x coincides with y at all times that x exists”¹⁶

¹⁴ I include this specification to allow for cases where the new donor organ, once transplanted, fails prematurely. In this case, Olivia would then need a second donor organ. For simplicity, assume that a given donor organ will last for as long as needed.

¹⁵ Presumably this is an unrealistic assumption. But it is a useful simplification for starting to develop the theory, since the complications which deterioration raises, as discussed earlier, are difficult.

¹⁶ I follow Hawthorne in this definition, as it one of the most neutral definitions available. Coincidence is defined in terms of overlap:

Overlap “ x *overlaps* y at a time just in case x and y share a part at that time”

I do not take a stand on which is more fundamental, the temporally extended object or its temporal parts.¹⁷ Rather, I point out this alternative because it reveals something about the nature of Olivia's claim to a temporally extended object in this case.

Suppose that we do elect to describe this object as many. We would then describe Olivia as having many claims, one to each of these objects. This is a useful way of conceptualizing Olivia's claims. She has a claim to the whole of the organ — the temporally extended object for that set period of time — as a result of having a claim to each of its temporal parts with that time period.

I said that this is useful, but the reader may wonder why. The answer is that we already do talk of claims to temporally extended objects in this manner, it just may not be obvious in the case of organs since they are more difficult to share between individuals. So, consider a different resource: a car. Suppose that Olivia owns a car, and then, at some time, t_{sale} , sells the car to her friend Sarah. Prior to t_{sale} , it would make sense to say that Olivia has a claim to the car (she owns it). But after t_{sale} , she no longer has a claim to it. Rather, Sarah now has a claim to the car. As noted above, this way of stating claims can be confusing: do we mean to say that Sarah has a claim after time t_{sale} to the car, or that, tenselessly, Sarah has a claim to those parts of the car after time t_{sale} . The latter, I suggest, is more clear even though not everyone will find this as intuitive. This is because it is *not* the case that at times after t_{sale} Sarah has a claim to the temporally extended object which is the car, for the parts of the car that existed prior to t_{sale} are not parts to which she has a claim. Her claim is *only* to those parts of the car that exist from t_{sale} forward (until such a time as she sells the car). Indeed, there are many resources for which the claims work in this manner. Another example is a house: individuals

Coincidence “x *coincides* with y at a time just in case x overlaps everything that y overlaps at that time and vice versa”

(Hawthorne, 2006, 85-86)

¹⁷ For a discussion of that debate, see Sider (2001), also Hawthorne (2006).

are not thought to have claims to their houses after they have passed away; rather, the later (temporal) parts of the house may be claimed by someone else.

The examples just given were in terms of property claims for ease of explanation, but the featurediscussed is also true of need claims. Suppose that, rather than one of Olivia or Sarah owning the car, they have a mutual ownership. Olivia needs the car at some time before t , and Sarah needs it for some time after t . In this case, it is true, tenselessly, that both have claims to the car — but to different temporal parts of the car.

These two results are useful, and can be combined to reveal something interesting about Olivia's claims to the organ: Olivia's claims to the later temporal parts of the organ *depend* on whether her claims to the earlier parts were satisfied. For, if those claims to the earlier parts are not satisfied, then Olivia will no longer be alive. As a result, those later temporal parts of the organ will not be instrumentally useful in satisfying her claim to D_A , and so, will no longer be needed. Since the basis of Olivia's claim was her need, if she no longer needs the organ's later temporal parts, then she will not have a claim to them. One might say that it is not until — or unless — her claims to the earlier temporal parts are satisfied that she might have claims to the organ's later temporal parts.

Might the same be true of Olivia's claims to the earlier temporal parts (just after $t_{\text{organ failure}}$), where Olivia would only have a claim to those earlier temporal parts of the organ if her claims to the later parts of the organ (up to and including t_{D_A}) will be satisfied? The answer depends on whether only having the organ for some of the time (as long as it is continuous with $t_{\text{organ failure}}$) can allow her claim to D_A to be partially satisfied. If it can, then it would seem that Olivia could have claims to the earlier parts of the organ even if her claims stop being satisfied at some time prior to t_{D_A} . If not, then Olivia would not have a claim to the earlier parts of the organ.

Is this a problem? If so, it is a red herring. Consider a case where an organ becomes available prior to $t_{\text{organ failure}}$. Above, I suggested that Olivia might have a claim to this organ. But it seems her claim to that organ depends on whether her claim to it at $t_{\text{organ failure}}$ (and some continuous times after that) will be satisfied. For if her later claim will not be satisfied, then it would seem that there is no reason for her to have this earlier claim: she does not need that organ (or those temporal parts of that organ), for it is not a means of obtaining the temporal parts of the organ which she would need.

7.2.1 MULTIPLE AVAILABLE ORGANS

Suppose, next, that there are *two* organs available in time for $t_{\text{organ failure}}$. Does Olivia have a claim to both?

Consider first the case in which both organs, organ A and organ B, become available at the same time. It does not matter which organ Olivia receives, so long as she receives one of the two organs, because she only needs one organ to satisfy her claim (provided the organ has the needed temporal qualities). So Olivia needs one of the organs, but it does not matter which. What does she have a claim to?

One might say that Olivia has a claim to both organs until such a time as she receives one of them. But something about this seems strange: Olivia would then have a claim to organ A, and a claim to organ B, but then either both claims go unsatisfied or one of the claims is satisfied and the other is neither satisfied nor dissatisfied. The former, both claims going unsatisfied, is not strange.¹⁸ But the latter seems bizarre: what happens to the other claim? Does it just disappear? It cannot be that the claim had only been to the use of organ B at times prior to

¹⁸ Indeed, such a description might have explanatory power regarding the intuition that those who have been waiting longer have a stronger claim than those who have not — they have more claims which have gone unsatisfied, and, in the interest of fairness, should have a claim satisfied before someone who has not yet had so many unsatisfied claims.

receipt of organ A, for the organ was needed as a means to having an organ during the needed times from $t_{\text{organ failure}}$ to t_{D_A} , a necessary condition for achieving lifetime well-being D_A . So that must mean that the claim disappeared.

Perhaps, the appearance and disappearance of claims should not be alarming. This seems to be a natural way of talking about what happens in other cases where property is transferred, or an individual no longer needs an object. But the reason why it is strange in this case is because the claim was never satisfied nor denied. In the other cases, the claims may usually be said to be satisfied or denied, since the claims were to objects at times, and when those claims “disappear” we can alternatively describe it as a case where there is no claim to those temporal parts of the object, as in the case of the car or houses, or to later parts of an organ that is of no use. And as was just noted, the disappearance of the claim to organ B (on receipt of organ A) cannot be redescribed in this manner.¹⁹

The problem may have to do with the reasons that the claims provide. To say that Olivia has a claim to both organs is to say that there is a reason for her to receive organ A and a reason for her to receive organ B. But these reasons are not independent of each other, since Olivia only needs to have one of the two.

To capture the dependence of these claims on each other, one might say instead that Olivia either has a claim to organ A, or that she has a claim to organ B, but not both. However, this analysis won’t work either: Suppose that Olivia’s claim is to organ A (and not to organ B). Suppose next that she does not get organ A — it

¹⁹ One might object: What if Olivia’s heart condition magically improved? Then surely her claim to an organ would disappear, and there would be nothing strange about this.

On the theory as presented, given the metaphysical assumptions noted above, this is an epistemic problem. Since, given the metaphysical assumptions, it would always be the case that Olivia’s condition would magically improve, neither organ would be necessary for her to achieve D_A . It may have *seemed* like she had a claim which then disappeared, but she, in fact, never had one.

This highlights the importance of the metaphysical assumptions taken. Unfortunately, I must save exploration of the influence of the metaphysics on our normative theories — and what we should do, given uncertainty about which metaphysics is correct — for another time.

is destroyed, or given to someone else, say. Would she then have a claim to organ B? The analysis under consideration already seems to assume that she cannot, since she cannot have a claim to both organs. Nor can we say that she had a claim to the earlier temporal parts of organ A, and that, when that claim was denied, she gained a claim to the later temporal parts of organ B: for, it might be the case that organ B is committed elsewhere — since it was unclaimed — at the same time as organ A.

A proponent of this analysis, that Olivia either has a claim to organ A or a claim to organ B but not claims to both, might then argue that in this case, her claim would always have been to organ B. For, if it was going to be the case that organ A went elsewhere, then organ A would not have been of use to her, and it could not have been the case that she had a claim to it. But there are two problems with this response. The first is that it misses the point: a claim is not a guarantee that the individual will get that to which they have a claim, nor should it be dependent on whether or not that particular claim will be satisfied. So whether Olivia has a claim to organ A is *not* dependent on whether her claim to organ A will be satisfied.²⁰ The second problem with the proponent's response is that Olivia might receive neither organ A nor organ B. In this case, it would not seem to make sense to say that, in light of organ A going elsewhere, Olivia's claim had always been to organ B. For then, might we not also say that, in light of organ B going elsewhere, Olivia's claim had always been to organ A? In such a way, it is unclear on this view what the claim was that was denied Olivia or the denial of which organ resulted in the denial of her claim, when it seems practically obvious that she has been denied her claim of need by being denied both organ A and organ B.

Ultimately, what this approach fails to account for is that both organs have the potential to allow Olivia to achieve D_A .

²⁰ Note that this is *not* the same as saying that whether an individual has a claim to one object depends on whether a *different* claim of theirs to a *different* object is satisfied.

What options remain, then, for describing Olivia's claim? I suggest that the best option is to say that Olivia has a single claim, to organ A or organ B. This is different from the last two options, because it does not assume that Olivia has more than one derivative claim (at least not more than one claim relative to the temporal parts of resources at a given time). The first option might have been written: (Olivia has a claim to organ A) AND (Olivia has a claim to organ B). The second option, similarly, might have been written: (Olivia has a claim to organ A) OR (Olivia has a claim to organ B) AND NOT ((Olivia has a claim to organ A) AND (Olivia has a claim to organ B)). This third option is different. It might be written: Olivia has a claim to (organ A OR organ B). Olivia's (singular) derivative claim, then, is satisfied just in case she receives organ A or organ B.

Why do I suggest that this is the best option? First, this way of describing Olivia's claim entails that her claim could be satisfied by the receipt of either organ, without a need to receive the other, and without the mysterious "disappearance" of any claims, since as soon as one organ has been received the claim is satisfied. Second, the claim is denied not by the denial of just one of the organs, but by the denial of both organs. This resolves the problems created by both of the other ways of describing claims.²¹

Still, one might object: why not abandon the project of claims to particular existent objects (or the disjunction of particular existent objects) in favor of a descriptive, of-the-sort-claim? This is the problem to which I turn next.

²¹ I gather that this is something like what Wiggins may have had in mind when he wrote:

[...] we may find it useful to be able to say that *y*'s need for *x* is *substitutable with respect to x* if some slight lowering of the standard by which *y*'s harm is judged permits us to weaken claims of need by disjoining *y*'s having *x* with his having *u* or *v* or *w* or whatever. Wiggins, 1998, 16

7.2.2 CLAIMS TO PARTICULAR OBJECTS OR TO DESCRIPTIVE (AN OF-THE-SORT CLAIM)

Why must derivative claims be to *particular* objects, rather than to a description? Why not say that Olivia has a claim to *an organ* with the needed characteristics (e.g. type of organ, time, etc.) in virtue of the necessity of such an organ for Olivia to achieve the level of well-being to which she has a primary claim? In such a way, Olivia's claim is not to any particular organ, but to getting *an* organ that matches this description.

This view might seem particularly appealing when considering claims to a share of something, be it a share of profits, harvest, or mushroom. Surely, one might argue, it does not matter which *particular* share the claimant gets — their claim is merely to a share of the right size. A proponent of this view might argue that this is a clear case where the individual's claim is not to a particular resource but to a resource which satisfies the right description. And their claim is denied just in case they do not receive anything of this description.

There are two problems with this view that claims are to *an* object of the sort whose description is given by the claim basis. The first is that this view cannot explain an analysis of a case where particular claimed objects have been destroyed. The second, related to the first, is that a step still needs to be made from a claim to *something* which meets the description to claims to *particular* things, rendering the claim to a descriptive an unnecessary step.

Suppose, for sake of argument, that this view is correct. Olivia then has a claim to an organ with the relevant features — to an “organ of description X”, where “X” is a placeholder for whatever those relevant features are. Suppose further that there are no other patients like Olivia who are in need of an organ of description X.

Suppose, then, that organ A becomes available. This is an organ which satisfies description X. But before Olivia receives organ A, Evilivia destroys it. Has Evilivia destroyed that to which Olivia has a claim?

Intuitively, the answer is yes. Olivia needed an organ of description X, she had a claim to getting *an* organ of description X, organ A — an organ of description X — became available, and so, it would seem, Olivia had a claim to that organ. Evilivia, in destroying organ A, destroyed that to which Olivia has a claim.

But the proponent of the view in question — that claims are to descriptives — is unable to say this just yet. For, on their view, Olivia's claim is to receiving an organ of description X, and not any particular organ. So, Evilivia might deny Olivia's claim to such an organ by destroying organ A, but she does not destroy *that to which* Olivia had a claim since Olivia's claim was not to any particular organ. But Evilivia's actions cannot even be said to have done this much if another organ of description X becomes available: if one does, call it organ B, then Olivia's claim will be satisfied if she receives it (and so Evilivia will not have denied her claim to an organ), or Olivia's claim won't be satisfied if she does not receive organ B, but then Evilivia can — at most — be said to have partially prevented the satisfaction of Olivia's claim.

Suppose that the proponent of this view wants to avoid this result. That is, they want to avoid this first problem by supporting the *prima facie* intuition that Evilivia has, in fact, destroyed something to which Olivia had a claim. They might then try to tell a story about how a claim to “something of X description” generates a derivative claim to a particular object that meets this description X.

This brings us to the second problem with this view, that the step from a claim to “something of description X” to a claim to “a particular thing of description X” renders the claim to the descriptive redundant.

Suppose that one did have a claim to *something* that meets description X. Then the derivation of a claim, from the primary claim to well-being to the derivative claim to a particular object would run as follows:

1. Olivia has a (primary) claim to well-being D_A in virtue of being the sort of agent she is. (*Assumed*)
2. To have a lifetime well-being of D_A , Olivia must live until t_{D_A} . (*Assumed*)
3. Olivia's (own) organ will fail at $t_{\text{organ failure}}$. (*Assumed*)
4. If Olivia does not receive an organ of description X by $t_{\text{organ failure}}$, then she will die at $t_{\text{organ failure}}$. (*Assumed*)
5. t_{D_A} is after $t_{\text{organ failure}}$. (*Assumed*)
6. If Olivia does not receive an organ of description X by $t_{\text{organ failure}}$, then she will not get lifetime well-being D_A . (*From 2, 4, 5*)
7. An agent has a derivative claim to a resource *of description* X, in virtue of a primary claim *iff* it is necessary that if the agent's primary claim is satisfied, then the agent receives a resource *of description* X. (*Modified Derivative Claim Principle*)
8. Olivia has a derivative need claim to an organ of description X. (*From 1, 6, and 7*)
9. There is an available organ, organ A, that meets description X. (*Assumed*)
10. An agent has a derivative claim to a particular resource Y *iff* she has a claim to a resource of description X and resource Y is a resource of that description. (*Bridging Principle*)
11. Olivia has a derivative claim to organ A. (*From 8, 9, and 10*)

But why add this redundant step? Olivia's claim to a particular object is not because it meets some description, but in virtue of the potential for *that particular object* to allow for the satisfaction of her primary claim. Why not just derive this claim directly from the primary claim to well-being? Since this latter type of derivative claim, to a particular object, still needs to be made, and could be done directly from a primary claim to well-being, including a claim to something that meets a description is unnecessary.

Still, it might be worth reviewing one of the reasons that the descriptive claim had seemed appealing in the first place: when an individual has a claim to a *share* of something, it would not seem that their claim is to a particular resource, but to a resource of the right description (here, size). The descriptive claim handles this nicely; but so too could the method discussed earlier for dealing with multiple objects (7.2.1). In cases where the claim is to a share of the profits, or the piece of mushroom, the individual in question has a claim to (slice 1 OR slice 2 OR slice 3 ...). Indeed, this might be a better description of the case. If an individual has a claim to a share of the mushroom of size X, then they have a claim to one of the particular slices of *that* mushroom. But suppose we write their claim as to "a slice of mushroom of size X" (i.e. a piece of mushroom of description X). Then, this claim will be satisfied provided they get a piece of mushroom of that size — regardless of whether the piece comes from *that* mushroom. Of course, one might alter the description to say "a slice of mushroom of size X from *that* mushroom"; but this is to switch from a claim to something that satisfies the description, to something that really is of the particular form suggested earlier.

7.2.3 EXISTENCE OF THE OBJECT NEEDED

This brings us to what is, perhaps, the most important condition on derivative claims of need: there cannot be a claim to a non-existent object. The reason for

its importance will be explained, but first, allow me to explain why this condition holds.

Suppose that Olivia lives in an unfortunate history where no organ ever becomes available. Her claim to a lifetime well-being of D_A has been denied. But is there any further claim that has been denied?

It seems that the answer must be no: there is no particular organ to which Olivia had a claim, since none ever existed. There was no particular resource to which a claim might attach. Derivative claims to resources, as just discussed, attach to *particular resources*. Had there been an organ, Olivia's claim to it would not have been in virtue of a claim to *an organ* (with the relevant features), but in virtue of the potential for *that particular organ* to satisfy her claim to well-being. This is similar to how Olivia did not have a claim to the organ that was available too late.

Perhaps, there is something about this that might seem strange. On the toy theory, Olivia's claim to a certain level of well-being has been denied, since she has not received that to which she had a claim. But well-being, like potential (but non-existent) organs, seems abstract. How could her claim attach to well-being in this manner if derivative claims cannot similarly attach to potential resources (rather than actual resources)?

I'm afraid I do not have a good answer to this question. The best I can offer is this: Suppose that derivative claims to resources could also attach to potential resources. How many potential resources might Olivia then have a claim to which would have been denied? Organs from donors, organs developed from animal parts, stolen organs? Organs available early, and available just in the nick of time? This would seem absurd.

Unfortunately, it is not a good response. The same response available in the case of multiple organs is available here: Olivia has a claim to a disjunction of po-

tential organs. There are not multiple claims to potential organs which have been denied, but rather one claim to a disjunction of particular potential organs which was denied (in virtue of none of those organs being actually obtained by Olivia).²² ²³

However, asking a different question might help provide motivation to the view that one can only have derivative claims to resources which exist, even if a clear answer to the question of why derivative claims cannot attach to potential resources is not forthcoming. The different question is this: is the denial to Olivia of the well-being really bad, in a way above and beyond the amount by which she is worse off? It might seem strange to say so: how could someone be denied a *claim* to well-being just in virtue of natural circumstances? On the other hand, however, this is not a particularly large stretch. “Life is unfair” is a cliché for a reason. The badness of living a life at a level below D_A is not just the difference in well-being between a life at the claimed level and the actual life, but also the additional bad of being denied a life that one deserves. Moreover, it seems a small price for a theory to pay if it helps explain the source and strength of claims of need to resources.

But now ask the same question again: is there anything else bad about Olivia not receiving an organ in the case where there is no organ to be had, above and beyond the amount by which Olivia is worse off and the badness of the denial of her claim to D_A ? It is difficult to see what. The denial of her claim to D_A is precisely

²² One might wonder: what if there were actual organs available? Would Olivia then have claims to both the actual and the potential organs? Surely, that would be absurd.

But again, the same escape is possible: if Olivia has one derivative claim, which is a disjunction of potential organs and the actual organs, then there is still only one derivative claim which is denied in the event she does not receive an organ.

²³What if her claim to actual organs is stronger than her claim to potential ones? This might explain why it would be worse for her to be denied an actual organ than to merely not receive one because it was only ever a potential organ. But it seems that if her claim to some of the organs would be stronger, it would need to be for some other reason. For example, if she was promised a particular organ, her (total) claim to this organ is strongest, but it is in virtue of having a claim of convention to it in addition to having a derivative desert claim to it. So, if a claim may take a disjunction of actual *and* potential organs, I’m not sure this distinction in strength can be made between claims to potential and claims to actual organs.

what would ground any need claim to an organ in the first place — that is the unfairness of not receiving an organ. But there is no further unfairness than this, since there is no particular organ which she could have had but has been denied.

Compare this to the case, discussed earlier in brief, where there is such a donor organ, available from some time prior to $t_{\text{organ failure}}$ and which will last until some time after t_{D_A} . Suppose also that there are no other patients like Olivia who are in need of this organ. The toy theory of claims, as I've presented it thus far, would assign Olivia (and only Olivia) a need claim to *this particular organ*. But the history takes a sad turn of events: though the organ exists at a time at which Olivia needs the organ, Olivia is denied the organ by her evil twin, Evilivia. Shortly after the organ becomes available, but before Olivia is able to secure the transplant, Evilivia sneaks into the hospital and destroys the organ for no other reason than sheer malice. Has Evilivia merely made Olivia worse off? No — Evilivia has also denied Olivia her claims, both to a certain standard of living, and to the particular organ that would have afforded her that standard of living. In this case, there is a further bad above and beyond the amount by which Olivia is worse off and the denial of her claim to D_A : it is the denial of her claim to *that* particular organ. She had a derivative need claim to it, and that claim was denied. This constitutes a further unfairness.

If this line of argument is correct, there are a few interesting conclusions which emerge. But before turning to those, there is an objection which must be met.

AN OBJECTION: PREVENTION OF THE EXISTENCE OF A NEEDED RESOURCE

Suppose that a technology has been developed which can produce organs of the sort Olivia needs by fermenting gray goop in large vats. The procurement of both gray goop and the vats is not difficult, and the fermentation process has started. Before an organ can be produced, however, Evilivia destroys the vats, thereby preventing

the production of the organ Olivia needs. *Prima facie*, Evilivia has done more than simply deny Olivia's claim to well-being; surely, Evilivia has also denied Olivia an organ.

On this theory of claims, however, this is not the case. Since Olivia can only have derivative claims of need to resources that *exist*, Evilivia didn't deny Olivia's claim to the organ — Evilivia prevented such a claim from even existing by preventing the existence of an organ. As a result, the only claim that Evilivia has denied — on this theory — was to the well-being, or so the objection goes. And if it holds, it appears to be a rather serious objection.

However, a similar response may be offered as before.²⁴ Had the organ production technology not existed, there would be no organ from it for Olivia to have a claim to, but we would not say in this case that her claim had been denied. Similarly, if an earthquake destroyed the vat, or killed the scientist who knew how to run the process, we would not say that her claim to an organ had been denied. What is the difference?

The difference appears to be agent interference. And I argued earlier that the claims we have should be protected, to a certain extent, from agent interference. But if Olivia can't have a claim to the organ that would have been produced, because its production has been prevented, how do we account for this?

The answer is actually relatively straightforward: Evilivia *did* destroy something that Olivia had a claim to, namely, the production process. The vats and gray goop are just as much to be considered resources as an organ. Indeed, given the manner in which derivative need claims work, it would seem that Olivia would have a claim to those resources being used in such a manner: the vats and gray goop are resources which are necessary for the production of the organ which is necessary for Olivia

²⁴ 6.3.1.

obtaining the level of well-being to which she has a primary claim.

This will become more clear if we suppose, contrary to our suppositions earlier in this chapter, that the production technology could also be used to produce kidneys, lungs, and livers as well as hearts. Olivia, who needs a new heart, couldn't have a competing claim with someone who needed a kidney, lung or liver — except insofar as they have a competing claim to the use of this technology for the production of the organ which they need.

In such a way, it seems, claims may backtrack. Olivia has a derivative claim, in virtue of her primary claim to well-being, to those particular, existent resources which will enable her to satisfy her primary claim. In the case just considered, this extends not only to an organ (if there is one), but also to the means for producing an organ.

7.3 TIME AND EXISTENCE OF THE CLAIMANT

The backtracking of derivative claims to resources raises two important questions, however. First, how far back in time can a backtracking derivative claim extend? This will be a difficult question, especially for purposes of determining the urgency of claims. It is evident that they can at least backtrack in time to a certain extent, as was made evident by Olivia's claim to the organ-producing vat and gray goop. But this raises the second — and for our purposes, more important — question: might claims of need backtrack so far that an agent has a claim to a resource at a time if that agent does not exist at that time? In particular, may an agent have a claim to a resource at a time *before* those times at which the agent exists?

I am afraid that I do not have a good answer to this question. But whichever way it is answered, it yields interesting answers to one of the questions I asked earlier about how time is relevant to claims of need. I'll discuss the implications of each answer in turn.

7.3.1 YES: CLAIMS TO RESOURCES AT TIMES *Before* THE TIME OF EXISTENCE OF THE CLAIMANT

Suppose that an agent may have claims to resources at times (or the temporal parts of resources at times) prior to that agent's existence. This doesn't seem like too far a stretch: those individuals would have a claim to a baseline lifetime well-being, like D_A , just in virtue of being the kind of agent that they would be. Plausibly, resources that exist before they are born might be necessary for enabling them to attain such a level of well-being. And, crucially, the non-identity problem would not seem to loom large: the claim of these individuals to the resources is not in virtue of avoidance of harm. Rather, their primary claim, to a baseline of well-being, is in virtue of being the relevant sort of agent. The claims to the earlier resources would then be in virtue of the necessity of those resources to their well-being, either by having those resources available at the later times or by having access to the products of those resources. As a result, it might not be known which agents will exist, but whichever ones do would have a claim, since this type of claim does not depend on counterfactuals in which those individuals may or may not exist.

One implication of this view is a temporal asymmetry with respect to those temporally located resources which might be claimed: derivative claims of need to resources might plausibly backtrack to times prior to the existence of the claimant, but are less likely to attach to resources located after the time of existence of the claimant. This was certainly the case for Olivia: she did not have claims to those temporal parts of organs which existed after her death because those organs would not be instrumental to her achieving that to which she had a primary claim. By contrast, should there have been an organ-producing technology available prior to her birth, she might have had a claim to it in virtue of its use in achieving the level of lifetime well-being to which she would have a claim. Provided that it is not a problem for agents to have claims to resources at times when the agent does not

exist, then Olivia would have a claim of this sort.

This result will turn out to be important, since it affects the manner in which the claims of the members of different, non-overlapping generations compete with each other. I will return to this in 9.1.2.

DOES BACKTRACKING WEAKEN A CLAIM?

If derivative claims can backtrack to times prior to the claimant's existence, a question arises about the strength of such claims: are they as strong as derivative claims which do not backtrack? It is tempting to think that they are not.

To aid in the discussion of this problem, I need to distinguish between direct derivative claims to resources and backtracked derivative claims. A *direct derivative claim* is one to a resource that *directly* contributes to the satisfaction of a primary claim. In our discussions, such a resource would directly contribute to the claimant's well-being. *Backtracked derivative claims*, on the other hand, are to resources that *indirectly* contribute to the satisfaction of a primary claim. The satisfaction of a backtracked derivative claim does not directly contribute to the satisfaction of a primary claim, but rather contributes to the satisfaction of the primary claim through a chain of such resources. For example, Olivia's derivative claim to organ producing technology (above 7.2.3) was a backtracked derivative claim: the organ producing technology could not directly affect Olivia's well-being, but rather provided another resource — an organ — which would directly contribute to Olivia's well-being. Call the number of steps in the (hypothetical) causal chain between a resource that indirectly aides in the satisfaction of a primary claim and the resource that would directly contribute to the satisfaction of the primary claim the *degree of backtracking*.

One reason why backtracked claims might be thought to be weaker than claims which are not backtracked is that the causal chain weakens as the degree of back-

tracking increases. The strength of a causal link between two events is in some sense dependent on the sensitivity to external circumstances of the nonoccurrence of the effect given the occurrence of the cause. The sensitivity of a causal chain to external circumstances seems to increase more than linearly with respect to the sensitivity of any of its links (e.g. probabilities multiply).²⁵ As a result, the longer the causal chain, the greater the insensitivity required of its links to be as strong a case of causation as a less insensitive causal link.

A second, and perhaps better, reason why claims with a higher degree of backtracking than those without has to do with two concepts introduced by Wiggins.²⁶ The first concept is that of *entrenchment*. The entrenchment of a need for a *particular* object is related to the substitutability of that object. It may be the case that the primary claim can be met through other means, but doing so will require some maneuvering. Wiggins writes:

... *y*'s need for *x* is *entrenched* if the question of whether *y* can remain unharmed without having *x* is rather insensitive to the placing of the aforementioned threshold of realistic envisage-ability-cum-political and moral acceptability of alternative futures. When we are concerned with the problem of arbitration between general needs claims or arbitration between general needs claims and other claims, it will then be important to distinguish between entrenchment with respect to the shorter term (where extant arrangements create definite requirements that cannot be escaped immediately but may be in due course escaped) and entrenchment with respect to the longer term. Some desirable disruptions of the established order that would enable people to escape harm without having *x* cannot be envisaged happening as it were overnight;

²⁵ Concerning the requirement of insensitivity, see Lewis (1986a, 184-188).

²⁶ Wiggins (1998, 14-16).

but in many cases the change can be coherently and easily described and can be quite realistically envisaged taking place gradually and by stages.²⁷

Entrenchment is related to a second concept introduced by Wiggins, that of *urgency*. The urgency of a need is how quickly a resource is needed to satisfy the primary claim. In many cases, it would seem that entrenchment and urgency go hand-in-hand: while the need for a particular resource might be substitutable if there were time enough to make appropriate changes, it might not be substitutable if the primary claim for which it is required is time-sensitive and soon since there would not be time to make alternate arrangements.

Resources with a higher degree of backtracking, it might seem, are less urgent in this regard, since the causal chain between such a resource and the resource that directly contributes is longer, and so, too, then is the temporal distance. This will be true of many, though not all, backtracked claims. Sometimes, the object so claimed will not be substitutable, even in the long term. But these are, in general, special cases.

I discussed above how to handle substitutable claims by use of a single claim to a disjunction of particular resources. It might be reasonable to conclude that the greater the number of particular objects within the disjunction, the less strong the claim is when in competition with other claims for one of those particular objects. If I have a claim to (organ A OR organ B), and you have a claim to (organ B) only, then it seems that your claim to organ B is stronger than mine — at least until such a time as organ A is no longer an option for me.

Working out the details of such claims compare in terms of strength based on the number of objects in the disjunction and on other details of the history in question would take us too far afield. For now, let us assume that the higher the degree of

²⁷ Wiggins (1998, 15).

substitutability — the greater the number of particular objects in the disjunction — the weaker the claim to any one of those particular objects.

Since the degree of backtracking seems to correlate in some measure with the degree of substitutability, it follows that, in most general cases, the higher the degree of backtracking, the weaker the claim.²⁸

There is a third reason to suspect that backtracking claims become weaker as they extend back in time prior to the existence of the claimant: if proximity is relevant to the strength of claims. I mentioned earlier in 5.5 that there are reasons to suspect this may be the case. Unfortunately, I lack the space to pursue this here.

For now, however, it seems reasonable to make the assumption that *if* claims can backtrack to resources at times prior to the existence of the claimant, that such claims become weaker with the degree of backtracking — and so too, then, with the temporal distance between the claimant and the object claimed.

7.3.2 NO: NO CLAIMS TO RESOURCES AT TIMES *Before* THE TIME OF EXISTENCE OF THE CLAIMANT

Alternatively, suppose that an agent may *not* have claims to resources at times (or to the temporal parts of resources at times) prior to that agent's existence. Indeed, this seems to be the standard line of thought when it comes to future generations.

²⁸ This seems to be related to a concern voiced by Parfit:

Suppose that we deplete some resource, but invent technology that will enable our successors, though they lack this resource, to have the same range of opportunities. There would be no objection to what we have done. The most that could be claimed is that people in each generation have a right to an equal range of opportunities, or to an equally high quality of life. (Parfit, 1984, 365)

Parfit and I differ in that Parfit uses this to argue that “people do not have rights to a share of a particular resource” (*ibid*). But, as I hope is clear from my discussion, this does not exclude the possibility that individuals might have *claims* to these resources. Rather, it demonstrates that the claims have a high degree of substitutability since the bases for such claims are not entrenched.

This view has a very important implication: the derivative claims to resources of members of different, non-overlapping generations will not compete. If the claimant must exist at the time of the temporal part of the resource claimed, then individuals will only have competing claims to resources with those other individuals who exist at the same time as their claim. For example, the members of future generations cannot be said to have claims to resources now — not because it is uncertain who will exist, or because such claims could not possibly be fulfilled²⁹ — but because they do not exist at this time. I will return to the implications of this for a theory of fairness in 9.1.3.

OBJECTION

Before closing, however, one might object that there is a serious problem with this view which combines the rejection of derivative claims to nonexistent objects with the rejection of derivative claims to objects at times when the claimant does not exist.

Recall the mushroom case presented in Chapter 6. Suppose that the members of the earlier generation consumed all or most of the initial mushroom, leaving little or none for the members of the later generations. According to a theory of claims which combines the two features just noted, the members of the earlier generation cannot be said to have violated the derivative claims of the later generations: there was no resource, to which those generations had a claim, which was destroyed. The members of the earlier generation affected what mushroom would be available for the members of the later generations to have a claim to, but this is not to deny a claim.

Is this a problem? It shouldn't be, at least not any more than any other theory which denies that later generations might have claims or rights which are what

²⁹ Parfit has a strange worry like this regarding rights to resources. *cf.* Parfit (1984, 364-366).

ground our obligations to them. Moreover, it doesn't mean that the earlier generation didn't have *any reasons* to conserve the resources for the later generations. Rather, it is just to say that claims to those resources are not among them. The satisfaction or dissatisfaction of claims are not the only things that matter on this theory. The value of well-being certainly provided reason to save or improve upon the available resources (since this has the chance to make things better overall). Similarly, the existence of a neutral level of well-being, above which it improves a history to add people, provides reason both to add more people to the history and to conserve resources to provide for them. I will return to this in 11.2.

7.3.3 DERIVATIVE DESERT CLAIMS TO RESOURCES AFTER TIME OF EXISTENCE OF CLAIMANT

After discussing the possibility of backtracking of derivative claims to times prior to the existence of the claimant, there arises a natural question about whether an individual may have derivative claims to resources (or the temporal parts of resources) after the time of that individual's existence. This was a question I had hoped to avoid, in part, by assuming *lifetime well-being*. But that very assumption also makes it somewhat unavoidable. Allow me to explain.

Prima facie, resources, or the temporal parts of resources, that exist *after* the time of existence of an agent would seem to not be instrumentally useful in increasing the well-being of that individual. Money, food, shelter — these resources are of no use after death. Nor might they be saved, preserved, or grown in a manner that will make them instrumentally useful in achieving well-being, unlike resources, or the temporal parts of resources, that exist *prior* to the time of existence of the individual. Recall that derivative claims to resources are in virtue of the object's potential for enabling the attainment of that to which the individual has a primary claim. Since it would seem that resources (or the temporal parts of resources) located after the

time of existence of the individual can be of no use in this way, an individual cannot have derivative claims to resources after the time of their existence. Indeed, this was the case with Olivia and the organ that came too late, discussed above.

But there are some views on which resources of the material sort can help increase an individual's well-being after they have passed, albeit indirectly. If the attainment of an individual's life projects after their death can increase their lifetime well-being, then it would seem that the resources that allow for the attainment of those projects have instrumental use in affecting the individual's well-being. If this is the case, then there does appear to be some scope for individuals to have derivative claims to resources located after the time of the claimant's existence.³⁰

There are two comments I wish to make about this possibility.

First, even where this is so, it seems reasonable to expect that resources are more likely to be of use during life than after it, since there are a greater number of ways in which those resources might contribute to an individual's well-being. In particular, the instrumental value of resources during life do not depend so heavily on an individual's projects, or their sustainment over long periods of time (including after that individual's death).

Second, would these claims be as strong? If such claims are indeed symmetric to the backtracking ones, then perhaps they would be. This will be the case where the reason why the backtracked claims are weaker is that the claimant is not "present" — that is, the claimant does not exist at the time of the claim. As I noted earlier, this view on the strength of claims at times when the claimant does not exist has *prima facie* plausibility. However, if the reason why the backtracked claim is weaker is due to the weakening causal chain or the availability of alternatives, then the

³⁰ I am here setting aside views that take material resources to affect an individual's well-being after death through the use of offerings. It seems that these views in some way assume that the individual still exists, just in a different form or realm. To avoid complication, I set this possibility aside.

symmetry is less clear. I noted that the reason why resources after an individual's existence might be instrumentally useful towards their well-being was that such resources might be required for the fulfillment of that individual's projects. In this way, derivative claims to resources after one's existence are not in virtue of a chain of "backtracked" claims.

Where does that leave us? Without a clear theory about how resources contribute to an individual's well-being, I am unable to make a decisive argument in favor of either position. But it does seem reasonable to assume, at least at first blush, that derivative claims to resources after the time of the claimant's existence will at least be less strong (if they are even possible). Given this, and the *prima facie* plausibility that resources after an individual's existence are not of use, I will assume that an individual may not have derivative claims to resources after the time of their existence. What follows from this assumption will ultimately depend on arguments for this position that must be left to another time.

7.4 CONCLUSION

This is as far as I will take our toy theory of claims. I return now to the problem of fairness and time, and the question of what our discoveries entail for comparative fairness.

Part III

The Problem of Fairness & Time,

Part II

In Part II, I started to develop a theory of claims that I hoped would provide support for the Progressive Scope Restriction (PSR) and the Temporal Comparison Thesis (TCT), thereby providing principled justification for Weighted Progressive Egalitarianism (WPE) (introduced in Part I). I also said that the plausibility of telic egalitarianism depended on the availability of such justification.

The project of Part III is to assess whether the results of Part II might provide this justification, and to evaluate the implications of the resulting theory, WPE. Accordingly, this part divides into several chapters. In the first, I at last explain the relevance and importance of claims to telic egalitarianism. In the second, I argue that the results from Chapter 7 do yield a telic egalitarianism which is not temporally neutral. The remainder of Part III is then dedicated to responding to objections (9.2), discussing considerations in formulating the new theory of Weighted Progressive Egalitarianism (Chapter 10), and suggesting potential applications for this theory that count in favor of developing it farther than I have been able to here (Chapter 11).

FAIRNESS, CLAIMS, AND THE DISVALUE OF INEQUALITY

In the introduction to Part II, I wrote that claims are important because claims are at the root of *fairness*, and the badness of inequality is that it is unfair.¹ The goal of this chapter is to explain how this is so.

The first step is to explain the relationship between claims and fairness. This is the goal of the first section. The second step, which is to explain how fairness contributes to the value of a distribution, is divided into two sections. The first of these discusses how complaints are measured: is it the inequality of well-being, objects claimed, or satisfaction of claims that contributes the disvalue? The second of these sections discusses how the relative strength of claims affects the size of complaints.

¹ *Supra* 81.

8.1 CLAIMS & FAIRNESS

When you have a claim to something, be it a resource or well-being, it is bad when it goes unsatisfied — and the person for whom it is bad is you. In such a way, the satisfaction of a claim is a personal good, or — in the language I have chosen to use — the *dissatisfaction* of a claim is a personal bad.

There is a further point to be made about this personal bad, which is that it is not yet comparative. In the case of Olivia, this was made explicit by assuming that she was the only person, and, more particularly, the only person with a claim to the organ. Regardless of who else existed in the population, or how those other individuals may have fared, if Olivia died a premature death, then she suffered the personal bad of not having her claim to well-being satisfied, and if there was an organ to which she had a derivative claim, then she also suffered the personal bad of not having her claim to that organ satisfied. Since she would suffer these bads — the dissatisfaction of her claims — regardless of how anyone else fared, the bads so suffered are noncomparative. What kind of bads are they? They might be subsumed under notions of justice or desert or something else. Or, perhaps, if fairness is about the value of claim satisfaction in any context, then it turns out that, contrary to what is generally thought, there is a noncomparative aspect to fairness.

It seems plausible to me to say that it is *unfair* for an individual's claims to go unsatisfied, and such a judgment is noncomparative. But I will not address here the appropriate categorization or label to put on these bads. Nor, unfortunately, will I address the more important question of how much disvalue the dissatisfaction of claims generates noncomparatively speaking. Finding the answers is difficult work left to another time.

Rather, my concern here is with a different set of questions about how the way in which the satisfaction of one person's claims compares relative to the satisfaction of another person's claims matters. That is, I am concerned with the disvalue

stemming from unfavorable comparisons of how the satisfaction of one individual's claims compare to the satisfaction of the claims of another. I take it that the standard view of what fairness is about has to do with a comparative theory of this sort.² This is because, on the standard view, fairness is about how individuals *compare* with respect to some relevant parameter — here, claims and the satisfaction thereof.

What, then, does fairness require? The answer to this question, when claims are equal, is relatively straightforward: fairness requires that *equal claims be satisfied equally*. When equal claims are not satisfied equally, this constitutes an unfairness. When an individual is worse off with respect to something claimed than another with an equal claim, this is an unfairness suffered by the worse off individual. As discussed earlier, the disvalue to the individual of this unfairness is called his *complaint*.³

But what fairness requires is less clear in situations where individuals have claims to different amounts of the relevant good (be it a resource or well-being), or claims of different strengths. One then might, following Broome, say that fairness requires that claims be satisfied in proportion to their strength.⁴ But there are also compelling arguments against the proportional view, like those provided by Kagan.⁵

I set these difficult questions about what fairness requires aside, and turn to a different one: how is the unfairness to be measured? For example, is it unfair if one individual has that to which they have a claim, but another who has the same claim has *more* than that to which they have claim? If so, how unfair? Should we measure an individual's complaint relative to that which he has a claim, or relative to what and how much those who also had a claim have? What about in cases where someone without a claim gets more than someone with a claim? Is fairness

² *e.g.* Broome (1984b,a, 1989); Kagan (ming, 434-435).

³ I follow Temkin in the use of this term (Temkin, 1993, 19-27). Broome similarly follows Temkin in this regard (Broome, 1989, 14).

⁴ Broome (1989, 13).

⁵ Kagan (2003).

only concerned with competing claims, or with the distribution of the object of those claims?

These questions boil down to questions about complaints, and how complaints relate to claims. This is the topic to which I turn next.

8.2 CLAIMS, FAIRNESS, & COMPLAINT BASES

My focus in this section is with what provides the complaint base — the relevant comparison that constitutes the unfairness and underlies an individual's complaint. I am not interested in fully developing a theory about how unfairness is to be measured; as has been demonstrated, entire books may be written on the subject which raise more questions than provide answers.⁶ Rather, I am interested in how primary and derivative claims affect *which* comparisons matter. To discuss this question, I return to a series of cases based on the case of Olivia and the organ from Chapter 7.

Recall that Olivia had a primary claim to a level of well-being, D_A , which could not be satisfied without an organ transplant. One such organ, organ A, becomes available. As a result of her need for the organ to achieve the well-being to which she has a claim, Olivia has a derivative claim to organ A. But I now make one change to the basic case: Olivia's claim to the organ is not the only one. There is another individual in this world, Sarah, whose plight is the same as Olivia's: she has a claim to a level of well-being, D_A , but needs an organ transplant in order to attain it. In such a way, Olivia's derivative claim to organ A is now *in competition* with Sarah's claim.

First, suppose that, before either Sarah or Olivia get organ A, the organ is destroyed. This is unfair in the noncomparative way noted earlier, in that both Sarah

⁶ Kagan and Temkin have done the field a great service in this respect. *cf.* Kagan (ming) (forthcoming); Temkin (1993).

and Olivia's primary and derivative claims are denied. But there is no *comparative* unfairness: since neither received that to which have a claim, their claims have been treated equally. As a result, there are no complaints. Call this Organ Case 1.

Consider, next, Organ Case 2, a variation on the original case. In Organ Case 2, Sarah gets organ A. As a result, Sarah's claims are satisfied, but Olivia's are not. Since Sarah and Olivia had equal claims, and their claims have not been treated equally, Olivia has suffered an unfairness. But how are we to measure the badness of this unfairness? In particular, what is the complaint base?

There are two claims with respect to which Sarah and Olivia have not been treated equally: the primary claim to well-being D_A , and the derivative claim to organ A. But there are two problems here. First, there is a dependency that exists between these claims, both with respect to the existence of these claims and their satisfaction. Second, there is a problem with scale. In particular, we need to get the organ on the same scale as the other stuff of value to Sarah and Olivia — we need a “currency” in which to value the organ. I'll discuss these problems in reverse.

What is the value of the organ, to either Olivia or Sarah? It would seem that the answer must be the value of the welfare it is able to provide. That is, the value of the organ is in enabling Olivia (or Sarah) to achieve the level of *well-being* to which she had a claim. Indeed, were the organ to enable not only the well-being which was claimed, D_A , but a lifetime with a higher level of well-being, greater than D_A , this seems to be lost value too. Whether the unfairness consists in the loss of claimed value, or loss of actual value, it seems that the appropriate way to measure the value of the claim must be somehow relative to the well-being it would provide.

Does this mean that the relevant comparison is between Olivia and Sarah's levels of well-being? It would seem so. For, as a result of the derivative claims not being satisfied equally, their primary claims to well-being are not satisfied equally. This is the unfairness done. As a result, the complaint base has something to do

with the comparison of well-being between Olivia and Sarah. This gives us a first weak principle about the disvalue of fairness: *a comparison of well-being between the well-being of two individuals matters when their derivative claims to the resources that allow for that well-being are not treated fairly.*

This still provides relatively little. Consider Organ Case 3. There is a third person in the world, Rob. Rob has a primary claim to the same level of well-being as Sarah and Olivia, D_A . He also has organ failure, and will not be able to achieve D_A without an organ of the same type as Sarah and Olivia. But Rob lives at a different time, and so needs an organ at a different time. There is such an organ, organ B, which would be of use to Rob in achieving D_A , but not of use to either Sarah or Olivia. Sarah and Olivia still have derivative claims to organ A, but Rob is the only one with a derivative organ claim to organ B.

As it turns out, Rob is lucky, and his claim to organ B is satisfied. Sarah is also lucky, and her claim to organ A is satisfied. Olivia has not been treated the same with respect to her claims. But what is the unfairness that she has suffered? That is, what is her complaint base?

As before, a comparison with Sarah seems relevant. Is a comparison with Rob also relevant? To answer this, we need to assess whether Olivia suffered an unfairness with respect to Rob. Is it unfair that Rob received organ B, but Olivia did not? No, since Olivia did not have a claim to organ B. So, is it still unfair that Rob attained well-being D_A , but Olivia did not? It is unclear how it could be: the well-being that Rob “received” was never Olivia’s to have. An individual’s claim to well-being is to a certain amount, not to a “particular lot” of well-being. So one cannot have a claim to the well-being of others (though one can have a derivative claim to the same resources in virtue of claims to well-being), so it would seem that no, Olivia does not suffer an unfairness in this respect either. As a result, the

comparison between Olivia and Rob is not morally relevant. Olivia has a complaint relative to Sarah, as before, but not relative to Rob. The addition of Rob to the case in this manner has not affected the disvalue of the inequality in this case, because Rob did not receive anything to which Olivia had a claim. This demonstrates a further condition on complaints: *an individual can only have a complaint against someone who gets that to which they had a claim.*

Some might interpret the result from Organ Case 3 as supporting the principle that complaints are only between those who have *competing* claims. But such a principle would be false. Consider Organ Case 4. There is another person, Evilivia, who lives at the time of Sarah and Olivia. She has the same claim to well-being D_A , but, since she does not need a donor organ to achieve it, she does not have a derivative claim to organ A. In this way, she does not have any *competing* claims with Sarah and Olivia. But, in Organ Case 4, organ A goes not to Sarah, but to Evilivia, who does not need an organ to achieve D_A and so who does not have a claim to the organ. As a result, neither Sarah nor Olivia's claims are satisfied. This is unfair: Evilivia got that to which they had a claim.

How do we measure this unfairness? That is, what forms the complaint base for Sarah and Olivia? One answer, as before, might be the difference between the well-being they actually get, and that to which they had a primary claim, D_A , which was the basis of their derivative claim to organ A. But what if Evilivia was fantastically well-off, and had far more well-being than that to which she has a claim. Is this even more unfair than if Evilivia is not fantastically well-off? Does it make a difference if receiving organ A made Evilivia better off (even though she did not need it to achieve the claimed level of well-being)? If one thinks the answer to these questions is yes, then the relevant comparison might be to Evilivia's well-being. I leave it open. The main point here is that the source of unfairness is not *competing* claims.

Rather, the source of unfairness is when another individual gets that to which you have a claim. When this happens, you have a complaint relative to them, regardless of whether they had a claim that competed with yours. *Complaints are relative to those who get those resources to which you had a claim.*

This gets us to one principle about the disvalue of inequality of well-being:

Disvalue of Inequality Inequality of well-being is bad when it reflects an unfair satisfaction of the derivative claims to resources.

This principle about the disvalue of inequality is based on this condition we just discovered about complaints:

Complaints and Claims An individual only has a complaint against those who have that to which the individual has a claim which has gone unsatisfied or been disproportionately satisfied (i.e. a claim in respect of which the individual has been treated unfairly).

This is an important result, but a full discussion of it must wait until Chapter 9. For now, there remains an important case that might be thought to be an objection to these principles.

Organ Case 5 is based on the original case with just Sarah and Olivia. There is a third person, Thea, who is a contemporary of Sarah and Olivia. She has the same claim to a baseline level of well-being, D_A , which is more than satisfied since she actually enjoys a level of well-being greater than D_A . As before, there is an available organ, organ A, to which Sarah and Olivia have derivative claims. Sarah receives organ A, and her claims to both the organ and to well-being are satisfied. It has already been shown that Olivia has a complaint relative to Sarah. But does she have one against Thea? Intuitively, it seems like she might: Thea has more well-being than that to which she has a claim, and Olivia has less. This seems unfair. But,

as described in this case, Thea does not have anything to which Olivia has a claim. As a result, Olivia cannot have a complaint against Thea. The intuition and the principle conflict.

This conflict might be a problem, but I do not think it is a serious one. Rather, I think the intuition might have a spurious source. In the case, it is stipulated that the only way for Olivia to achieve D_A is to get an organ. It is the lack of this resource that causes her to be worse off, but she is not in competition with Thea over it. But the reality is that there might be other ways to compensate Olivia and help her attain a higher level of well-being that approaches if not reaches D_A .⁷ The way to compensate would require the use of other resources to which she would have a claim. It is plausible to suspect that Thea might have at least some of these resources. Thea would have a claim to some of them; but there may be many to which she does not have a claim, at least not a derivative one since Thea is already better off than D_A . If this is the case, then there is an unfairness that Olivia suffers relative to Thea.

This objection, then, does not constitute an objection to the principles presented above, but rather highlights a complication that arises as a result of these principles. Unfortunately, the complication will be difficult to make sense of or resolve until we can arrive at a clearer theory of how resources contribute to well-being. But finding such a theory is work for another time.

Until such a time as we are able to develop that sort of theory, we might still want to say something useful about how different complex inequalities compare, and which comparisons matter. So, for the remainder of this thesis, I will make two simplifying assumptions: First, all the resources at the time at which an individual lives are instrumentally useful to that individual by increasing their well-being. Second, everyone is the same both in terms of their needs and of how those resources

⁷ I set this aside earlier for simplicity.

contribute to their well-being. The result of these simplifying assumptions is that if someone has more well-being than another person, it is because they have more resources during the time at which they live.

8.2.1 STRENGTH OF CLAIMS AND SIZE OF COMPLAINTS

Before turning to the problem of discussing how the results of Chapter 7 affect what we can say about the disvalue of inequality, there is one final question about how the relative strength of claims affects complaints. My goal here, again, is not to fully work out the answer to this question. Rather, it is to consider whether there are any principles which any answer would need to satisfy.

Suppose that, as before, Olivia and Sarah both have a primary claim to well-being D_A , but cannot achieve it without an organ of the appropriate sort. Suppose also, that there is again a third person in this world, Rob, who lives just before Olivia and Sarah, and, like Olivia and Sarah, has a primary claim to D_A that he cannot achieve without an organ. In addition, there is a fourth person, Evilivia, who lives at the time of Olivia and Sarah, also has a claim to D_A , but unlike the others, can achieve it without an organ.

Fortunately, there are three organs available: organ A, organ B, and organ C. Organ B is the only one available at the time when Rob lives, and so is the only organ to which Rob has a derivative claim. But, unlike in our previous story, organ B could be preserved and so would be useful to Olivia and Sarah in achieving D_A as well. As a result, both Olivia and Sarah have backtracking claims to organ B. Organ A and organ C are available at the time when Olivia and Sarah exist, and would also be useful for them to achieve D_A . Accordingly, Olivia and Sarah have a derivative claim to (organ A OR organ B OR organ C).

The time comes to distribute organ B. Whose claim to it is strongest? The

answer would seem obvious: Rob's is, since his claim is entrenched.⁸ Suppose that Rob does get organ B, and further, that organ A is given to Sarah. Unfortunately for Olivia, however, organ C goes to Evilivia. (Evilivia wants the organ to put on her mantel as a work of "modern art".)

Olivia suffers an unfairness, but what is the extent of it? She suffers an unfairness relative to all three other individuals in the history, since they each received that to which she had a claim but was denied. I set aside the problem of how to combine these unfairnesses to arrive at Olivia's complaint. Rather, I ask this question instead: does Olivia suffer the same unfairness relative to each of them?

It would seem the the greatest unfairness she has suffered is relative to Evilivia. This is because Evilivia did not even have claim to organ C. Put another way, any claim Evilivia had to organ C paled in comparison to that of Olivia's. Olivia's claim to organ C was far, far stronger than that of Evilivia.

This seems to suggest a rather natural principle: it is a greater unfairness for a claimed resource to go to an individual with a lesser claim than to an individual with an equal (or stronger) claim. Hence, complaints are greater — that is, complaints have greater disvalue — when the individual who suffers the unfairness had a stronger claim to the resource than the individual who received that resource.

We can again try this out relative to Rob and Sarah. Olivia suffers an unfairness relative to Sarah, since she had a claim of the same strength as Sarah to organ A. This unfairness, however, is less than the unfairness which Olivia suffers relative to Evilivia. Indeed, this seems plausible. Olivia also suffers an unfairness relative to Rob, but this seems like it is less unfair: Rob had a stronger claim to organ B than Olivia. This has at least *prima facie* plausibility. It may not be obviously true, but it also is not obviously incorrect. Which it is will depend, in some measure, on

⁸ *Supra* 166.

questions left open about competing claims that are disjunctive — on whether or not, in light of the turn of events with Evilivia, Rob *actually* had a stronger claim than Olivia to organ B. The source of doubt about the principle's application in this case is not whether the principle is correct, but whether Rob's claim to organ B was in fact stronger than Olivia's. So, it would seem, we have found a plausible condition on how the relative strength of claims affects complaints. This will turn out to be useful.

And so, at last, I return to the question of fairness and time.

FAIRNESS & TIME

This chapter, though not the concluding one, is the chapter that pulls together the results of the previous sections to argue in favor of a telic egalitarianism which is not temporally neutral. The first half of the chapter discusses the argument in favor of such a position. The second half of the chapter responds to objections.

9.1 THE ARGUMENT

There is a basic structure to the argument pulling together the results from the previous chapters. It depends, to some extent, on questions that were left unanswered. But, as I noted at those places, whichever way the questions were answered, there were implications for telic egalitarianism that might help it avoid the objection raised in Part I.

The basic structure to the argument pulling together the results from the previous chapters is as follows:

1. There's greater overlap of instrumentally useful resources between contemporaries.

2. Therefore, the claims of contemporaries are in greater competition than the claims of noncontemporaries.
3. Inequality of well-being is bad when it reflects an unfair satisfaction of the derivative claims to resources.
4. Therefore, the badness of inequality between contemporaries is greater than the badness of inequality between noncontemporaries.

This is a rough form of the argument, given so that the reader can see the general outline of how the different results come together. But it skips steps, and is stated imprecisely, in terms of inequality rather than complaints. It also overlooks the varying implications of the questions that remained unanswered.

In what follows, I present the actual series of arguments that yield this result using the results from previous sections, with variations based on the answers to the open questions. The chapter is divided accordingly. There were two questions left open: first, whether individuals could have derivative claims to resources at times before those individuals exist; and, second, whether the backtracking of a claim leads to a diminution in its strength.

Throughout, I make four simplifying assumptions, which I will refer to below as follows:

Assumption 9.1.1. Everyone is equally morally deserving of the same desert baseline D_A . (i.e. Everyone is equally deserving, with respect to both noncomparative and comparative desert.)

Assumption 9.1.2. All the resources at a time at which an individual lives are instrumentally useful to that individual: if they receive the resource, it will increase their well-being.

Assumption 9.1.3. The manner in which a resource affects an individual's well-being is the same for all individuals.

Assumption 9.1.4. There are no other resources in the world which affect an individual's well-being.

As a result of these assumptions, we get:

Assumption 9.1.5. The amount of well-being an individual has is proportional to the amount of resources that individual has.

9.1.1 THE ARGUMENT AGAINST COMPLAINTS RELATIVE TO THE FUTURE

The first argument makes use of the result that a claimant cannot have derivative desert claims (hereafter, “derivative claims”) to resources (or the temporal parts of resources) located at times after the existence of the claimant (7.3.3). An important disclaimer, however, is that I did not argue that this assumption was correct. Rather, I argued that, in the absence of a better theory about how resources contribute to an individual's well-being, it was plausible, and — more importantly — a good approximation of the general case.

The argument, then, is as follows:

1. An agent cannot have derivative desert claims to resources (or the temporal parts of resources) located at times after the existence of the claimant. (*From 7.3.3*)
2. An individual only has a complaint against those who have that to which the individual has a claim in respect of which the individual has been treated unfairly. (*Complaints and Claims, from 8.2*)
3. An individual cannot have a complaint against those individuals who exist entirely at times after which he exists (i.e. who have no temporal overlap with him *and* who live after the time at which he exists). (*From 1 and 2*)

The conclusion of this argument, Point 3, is the *Progressive Scope Restriction*. Views that satisfy this restriction are members of a family of views called *Progressive Egalitarianism*, which I introduced in 4.3.3.

9.1.2 THE ARGUMENT: YES, CLAIMS CAN BACKTRACK PRIOR TO CLAIMANT'S EXISTENCE

The main appeal of Progressive Egalitarianism was that it matched the *prima facie* intuition in response to the Simple Temporal Case (1) that (c*) was worse than (a*). For this to be hold, it must be the case that individuals may have complaints against earlier individuals with whom they do not have any temporal overlap. This, in turn, depends on the first question left open, namely, whether claims might backtrack to times prior to the claimant's existence.¹ In this section, I'll consider the implications if the answer to that question is yes.

The argument proceeds as follows:

1. Derivative desert claims may backtrack to resources at times prior to the existence of the claimant. (*From 7.3.1*)
2. An agent cannot have derivative desert claims to resources (or the temporal parts of resources) located at times after the existence of the claimant. (*From 7.3.3*)
3. An individual only has a complaint against those who have that to which the individual has a claim in respect of which the individual has been treated unfairly. (*Complaints and Claims, from 8.2*)
4. An individual may only have complaints against those with whom he has some temporal overlap or who live before him. (*From 1 and 3*)

¹7.3, *supra* 163.

Egalitarian theories that satisfy the condition set out in 4, above, form a subset of Progressive Egalitarianism (since they also satisfy the Progressive Scope Restriction). These theories stand well-placed to yield the *prima facie* judgments of the Simple Temporal Case. For them to do so, however, depends on two features: first, on how complaints work when satisfaction relative to the primary claim is uneven due to the distribution of the derivatively claimed resources, and, second, what levels of well-being were represented by the distribution. The difficulty is that there is still quite some way to go towards answering the question about how unfairness is measured when absolute desert has been violated owing to either a surplus or scarcity of derivatively claimed resources.

To demonstrate how such an argument might run, however, I'll lay out the steps for the argument from one set of the relevant assumptions that yields the conclusion that (c^*) is worse than (a^*) with respect to inequality. To do so, I'll continue the argument above by assuming that at least some people pictured in the distribution have a higher level of well-being than D_A . Given the assumptions set out above (9.1), this can only be a result of an uneven distribution of the available resources.

5. At least some people pictured in the distributions of the simple case have a higher level of well-being than D_A . (*Assumed*)
6. In (a^*) , no one who has a well-being higher than D_A is using resources to which another individual has a derivative claim. (*From Assumption 9.1.5 and 2*)
7. There are no complaints in (a^*) . (*From 3 and 6*)
8. In (c^*) , individuals who are worse off than D_A have derivative claims to resources located at earlier times. (*From 7.3.1 and Derivative Claim Principle, 5.3.2*)

9. In (c^*) , those who are better off than D_A are using resources that others, who are worse off than D_A , have a derivative claim to. (*From 9.1.5 and 8*)
10. In (c^*) , the worse off are being treated unfairly relative to the better off with respect to their claims. (*From 9*)
11. In (c^*) , those who are worse off than D_A have complaints against those who are better off than D_A . (*From 10 and 3*)
12. (c^*) has more complaints than (a^*) . (*From 7 and 11*)
13. (c^*) is worse than (a^*) with respect to inequality. (*From Definition of Complaints in 4.1 and 12*)

Whether this conclusion will result, irrespective of the assumptions about the level of well-being present in the Simple Temporal Case, depends on the answers to complicated questions about when an imbalance of well-being, owing to an uneven distribution of resources, is unfair. As already noted, this is a complicated topic, one which we unfortunately lack the space to explore further. What is interesting is that in at least one plausible description of the Simple Temporal Case, our earlier results, combined with a plausible answer to an open question, led to the appropriate judgment.

9.1.3 THE ARGUMENT: NO, CLAIMS CANNOT BACKTRACK PRIOR TO CLAIMANT'S EXISTENCE

The result of the last section depended on the potential for derivative claims to backtrack to times prior to the claimant's existence. However, the question of whether that is possible was left open (7.3). In this section, I will make the assumption that the answer to this question was that claims may *not* backtrack to times prior to the existence of the claimant.

The argument from this position, then, is as follows:

1. Derivative desert claims may not backtrack to resources at times prior to the existence of the claimant. (*From 7.3.2*)
2. An agent cannot have derivative desert claims to resources (or the temporal parts of resources) located at times after the existence of the claimant. (*From 7.3.3*)
3. An individual only has a complaint against those who have that to which the individual has a claim in respect of which the individual has been treated unfairly. (*Complaints and Claims, from 8.2*)
4. An individual may only have complaints against those with whom he has some temporal overlap. (*From 1, 2, and 3*)

Point 4 is a strong form of the scope restriction I suggested was plausible in 4.3.3. It also satisfies the Temporal Comparison Thesis, that inequalities between those who are temporally close matter more than those who are temporally distant, albeit, in an extreme manner.

Is this plausible? It certainly does not match the intuitions presented in the Simple Temporal Case (1). But perhaps those turn out to be incorrect. In 7.3.2, I raised and responded to the objection that the members of later generations should be able to have claims to resources that exist before them, for otherwise one cannot say that earlier generations who deplete resources have violated the claims of later generations. There I said that this should not be anymore troubling than any other theory which denies that later generations might have claims or rights which ground our obligations to them. Indeed, the growth models provide a clear reason why (c*) (from the Simple Temporal Case) is worse, all-things-considered, whereas (c) is not. The reason is this: if the growth models are correct, then the members of earlier generations who have depleted resources have made the history go much,

much worse. But this is not a matter of fairness.

This might appear to be an argument in favor of valuing equality-at-a-time. Since there are arguments against such a position,² this might appear to be a damaging objection to the assumption that claims to resources may not backtrack. However, the result of this argument is *not* that inequality is only bad *at-a-time*. Rather, it is the position that comparisons of well-being only matter between individuals who have at least some temporal overlap.

This can be shown by removing the assumption of non-overlapping generations. In this case, if an individual has less than D_A while someone with whom he has temporal overlap has well-being higher than D_A , then he has not been treated fairly with respect to his derivative claims. This is because, given our assumptions above, the better off individual has some of the resources to which the worse off individual has a derivative claim. This will be true regardless of when during the life of the worse-off individual the temporal overlap occurs. That this is so is a result of the assumption of lifetime well-being, and the even contribution of resources to well-being. Varying either of these two assumptions will vary the result. However, what the simplified version shows is that there can be unfairness arising between members of different generations provided that those generations have temporal overlap.

While it might be plausible that members of generations in the far future cannot have backtracking claims to our resources now, this seems less plausible in the case of the next generation, where it might well be the case that a worse off individual in the next generation has no temporal overlap with one of the very best off individuals in this generation owing to a mere twist of fate. That this position does not allow for *any* unfairness between members of consecutive generations who do not have temporal overlap provides some reason to doubt that derivative desert claims may

² e.g. Special relativity makes simultaneity arbitrary. (Thanks to Will Crouch for this point, which Will credits to Toby Ord.)

not backtrack prior to an individual's existence. For this reason, I will continue to explore developing a temporally-sensitive egalitarian theory on the assumption that derivative desert claims may backtrack to times prior to the existence of the claimant to enable further discussion.

9.1.4 THE ARGUMENT: CLAIMS WEAKENED BY BACKTRACKING

There arose earlier a question about how the backtracking of derivative desert claims affects the strength of those claims (7.3.1, page 165). Although I left it an open question, I suggested that it was plausible that the higher the degree of the backtrack, the weaker the claim. As just noted, I am going to assume that claims may backtrack to times prior to the existence of the claimant.

The argument in this case is as follows:

1. Derivative desert claims may backtrack to resources at times prior to the existence of the claimant. (*From 7.3.1*)
2. Derivative desert claims lose strength as they backtrack in time. (*From 7.3.1*)
3. An individual only has a complaint against those who have that to which the individual has a claim in respect of which the individual has been treated unfairly. (*Complaints and Claims, from 8.2*)
4. Complaints have less weight where the claims involved are less strong. (*From 8.2.1*)
5. An individual's complaints will be less strong when resulting from backtracked claims. (*From 2 and 4*)
6. An individual's complaints against those who exist prior to him will result from a backtracked claim. (*From 7.3.1*)

7. An individual's complaints against those who are temporally distant will have less weight than his complaints against those who are temporally close or contemporaneous. (*From 4, 5, and 6*)

Point 7 is the Temporal Comparison Thesis, the thesis that, where inequalities matter, inequalities between those who are temporally close matter more than inequalities between those who are temporally distant. We can now see the reason why: it is due to the weakening of backtracking claims — the only sort of claim that can attach to resources that are temporally distant from the claimant — which results in weaker complaints between individuals who are temporally distant.

When the TCT is combined with the Progressive Scope Restriction, the result is *Weighted Progressive Egalitarianism*. This is the position that I suggested might help save telic egalitarianism if justification for it could be found. Unfortunately, I have only been partially successful in this endeavor. I was only able to provide partial arguments or reasons to find plausible the underlying assumptions that lead to this position: the lack of derivative claims to resources at times after the existence of the claimant (7.3.3), the possibility of backtracking claims to resources at times prior to the existence of the claimant (7.3.1), and the weakening of backtracking claims (7.3.1). Still, this constitutes some progress: the position has been shown as plausible, and the points in need of further argumentation have been made clear. Ultimately, they rest on the development of a theory of how resources contribute to well-being, a task far beyond the scope of the present work.

In the remaining chapters, I explore the implications of this view for which I have argued. Chapter 10 explores various ways of formalizing WPE, and considerations that should be kept in mind when doing so. Chapter 11 explores some potential advantages of the theory which provide reason for taking the next steps towards demonstrating and developing WPE.

But first, however, I must answer a few objections.

9.2 OBJECTIONS

There are three sets of objections that might be aimed at the theory just argued for. The first constitute potential objections to the theory itself. If not met, these objections cast doubt on whether the view under consideration is even worth pursuing. The second set of objections regard whether Weighted Progressive Egalitarianism really constitutes a way of recovering telic egalitarianism. In this way, the second set is less serious, though it is of interest. The third set concerns whether this theory is about time, or something else. I conclude with a brief discussion of another possibility that has frequently been suggested in response to my proposed theory: that the TCT does not track temporal distance, but weak causal connections.

9.2.1 IMPARTIALITY & RELATIVITY: AN OBJECTION TO WEIGHTING COMPLAINTS?

One might worry that weighting complaints for temporal distance would violate impartiality, result in a relativistic theory, or both. Whether and why such violations or results should be avoided is best left to another discussion, but the trade-off is great enough that one would need a very compelling account in favor of any theory with these results, and a strong explanation of why the trade-off is warranted.

Fortunately, this is not the case with weighting *comparative complaints* for temporal distance.

First, weighting comparative complaints does not violate impartiality. It treats each individual the same: each individual has the same function for the contributive value of their well-being to the good. Permuting the *identity* of individuals — without affecting the *relative* “location” (spatial or temporal) of well-being — does not change the overall value of a particular outcome or history.

Note that this differs from a theory that involves temporal discounting of an

individual's well-being by discounting personal good. A theory that discounts in this manner *does* violate impartiality.³ For example, a theory that discounts the value of future well-being as compared to present well-being is partial to the present. The source of the confusion might be that, though permuting the identity of individuals does not, in itself, affect the overall value of a particular outcome or history, changing the temporal "location" of the well-being *does*. This may be the source of the confusion.

Second, a weighted theory of equality is not relativistic. The ranking — or value — assigned to each outcome does not change depending on viewpoint. Rather, the size of complaints depends, in part, on the temporal distance between the relevant parties.⁴ It does *not* say that those closer to you *matter more*; rather, it says the *comparison* does.⁵ The former is not plausible, but the latter might be.

9.2.2 IS WEIGHTED PROGRESSIVE EGALITARIANISM REALLY AN EGALITARIAN VIEW?

WPE IS NOT AN *Egalitarian* VIEW

One might object that WPE is not really an egalitarian view. There are several sources for this objection.

First, while WPE may satisfy some definitions offered in the literature, these definitions might be seen as misleading. For example, Parfit writes that the *Principle of Equality* is:

³ The main problem with such views is that they are not actually compatible with consequentialism, despite purporting to be consequentialist theories. See Cowen (1992).

⁴ It may also turn out to depend on spatial distance. I have not considered such a position here, but Weighted Progressive Egalitarianism is not incompatible with a view that would.

⁵ This may have the implication that, for certain distributions, to transfer a unit of well-being (or the relevant 'good') to someone temporally near is better than the same transfer to someone temporally far. This will be so in situations where such a transfer would reduce the close local inequality more than it would reduce a local inequality at the temporally distant location.

It is in itself bad if some people are worse off than others.⁶

WPE counts as an egalitarian under this view: it is bad if some are worse off than others, namely, when their claims have been treated unfairly. But Parfit also writes:

On these definitions, we are Egalitarians if, in any area, we believe we should aim for equality. If we had that belief in only some small area, we would not naturally be called ‘Egalitarians’. In that respect my definitions are misleading.⁷

Is this a problem for WPE? This is best answered by considering its two primary features.

The Temporal Comparison Thesis does not, on its face, violate egalitarianism: we have not yet denied that there is something (however weak) in favor of equality between temporally distant individuals. Rather, it is a question of degree: not all *comparisons* are made equal. Indeed, this would not be the first time it has been questioned whether the same weight should be given to all comparisons of inequality. There is much egalitarian debate over whether such comparisons matter more in rich societies or poor societies, for instance.⁸ If it is bad that someone is worse off than another through no fault of his own, there still remains the possibility that it can be more or less bad depending on where she is relative to where the other person is.

The Progressive Scope Restriction, however, might be more problematic, since there will not be a complaint if an earlier person is worse off than a later person. In this way, WPE does not care about inequalities when an earlier person is worse off than a later person. This does not entail, however, that WPE does not aim for equality, or that it only does so in “one small area.” WPE cares about equality between individuals when they are equal with respect to their claims, and this is a

⁶Parfit (1995, 84). He remarks that we might add the qualifier “through no fault or choice of theirs” in a note to this definition.

⁷Parfit (1995, 122 n.7).

⁸Temkin (1993, 157-190).

rather large area. It is also the case that a concern for equality can be made with caveats such as this one. As Blackburn points out, “in Aristotle equality of treatment is contingent upon the subjects of the treatment being equal ‘in the relevant respects.’ ” The difference between WPE and TNE, then, is over whether temporal location is a relevant respect.

A second source for this objection is that egalitarian views take the value of equality between individuals — or, in the case of deontic egalitarians, the aim of equality — as basic. Since WPE bases its theory on claims, it is not properly called “egalitarianism.” But WPE does take a sort of equality as basic, namely, equality with respect to the satisfaction of claims, so this objection does not seem well-founded.

Either way, I am not particularly interested in debating this point. My hope is that it constitutes the start of a plausible theory for distributive ethics. If it is not a “true” egalitarian theory, so much the worse for egalitarianism.

WPE IS NOT AN *Teleological* EGALITARIAN VIEW

One might object that WPE is not a *teleological* egalitarian view, since it takes inequality to be bad only when it reflects *unfair* satisfaction of claims. But this objection, again, is misplaced: WPE doesn’t simply say that *inequality is unfair*; it says that unfair inequalities are *bad*.

Again, however, I am not sure that much turns on this classification.

9.2.3 TCT IS NOT ABOUT TIME

The final objection is that the TCT does not really “discount” complaint bases for *time*. Rather, it discounts complaint bases for the weakness of claims owing to backtracking. It just so happens that there is a difference between the strength of

claims of temporally distant individuals to the same resource, but this difference is not due to time but to due to the backtracking.

This is partially correct. The “discounted complaint bases” arise because the claim of the worse off is weaker. But the reason that claim is weaker is owing to a comparison of the *temporal distance* between the two individuals and the temporally-located object so claimed. So the discount in the complaint base can reasonably be called a discount for time.⁹

CAUSATION & INEQUALITY

Before continuing, it is appropriate to pause and return to a related question about causation.¹⁰ Do temporally distant comparisons matter less because the causal connections are weaker?

Not exactly. Causal connections, as we have seen, do play a role. But they play a relatively minor one: it is not causal connections between *individuals*, but causal connections between individuals and temporally locatable *resources* that matter. Indeed, the weakening over distance of causal connections to resources was identified as a potential reason for the weakening of backtracking claims.¹¹

But it is also important to articulate the role that causation does *not* play. Inequalities relative to other people can be bad even if they are not *caused* by

⁹ Broome makes a similar point about discounting commodities:

Is the discount a discount for time? Parfit says no: since future commodities are discounted because people have more future commodities than present ones, it is misleading to say the discount is for time. But I do not think it is misleading. It is commodities we are discounting, not wellbeing. Dated commodities are identified by their dates, and it happens that future commodities are cheaper than present ones. This can reasonably be called a discount for time. What *causes* future commodities to be cheaper than present ones is another matter. Evening phone calls are discounted compared with daytime phone calls. The cause of this discount is that fewer people use the phone in the evening [...]. Nevertheless, the discount is for the time you make the call. It is evening calls that the phone company markets at a discount, not calls when few people are using the phone. (Broome, 1999, 59-60)

¹⁰ *Supra* 3.2.

¹¹ *Supra* 7.3.1.

the actions of those people, and even if those people are helpless to remove the inequality. That is, the view argued for does not reduce telic egalitarianism (the view that inequality is intrinsically bad) to deontic egalitarianism (the view that inequality is only bad insofar as it results from some unfair process).¹²

Consider a temporally local, or “static”, case of inequality without “unfair” causation — that is, a case where the inequality resulted through no fault of the individuals concerned:

Earthquake Inequality Case There exists a generation of individuals alike in all morally relevant respects who, for the first part of their lives, enjoy the same level of well-being. At some point partway through the lifetime of this generation, an earthquake strikes, rendering half the population worse off than the other.

Could this situation be improved by a mere redistribution of the suffering? It seems that it would be: it would be better if the same total loss of well-being were shared by many, with each losing a little, than served as a heavy burden shouldered by a few. This, I suggest, is part of the reason we have an obligation to aid in such cases — even when providing aid is not achieved by an efficient transfer. It also seems that the evenly distributed smaller burdens would be better even if they cannot be achieved, as happens when an earthquake inflicts an unrecoverable burden on an individual (by killing them or placing them beyond aid). Indeed, it seems that we have reason to try to find ways of getting aid to such individuals, by, for example, improving remote medical assistance technologies. That it would be better if we could assist them and spread the burden provides one reason to improve our technological capacities in this manner. This is an example where the disvalue of inequality provides reason to

¹² Parfit has characterized the latter view as: “What is unjust, and therefore bad, is not strictly the state of affairs, but the way in which it was produced.” (Parfit, 1995, 90)

If the proposed modification of telic egalitarianism is unsuccessful, this *reductio* may be the end result. An earlier version of Part I led one reviewer to remark that the problem raised seemed like a *reductio ad absurdum* of telic egalitarianism to deontic egalitarianism.

develop causal connections where none would be otherwise.¹³

Supposing that this is correct, there are three ways to capture this ranking of the alternatives in our axiology: through the intrinsic disvalue of comparative complaints, through a prioritarian view, or by treating uneven distributions as a communal bad. I earlier set aside the latter two views.¹⁴ Insofar as the first method for capturing this ranking is the correct one to apply, then the Earthquake example shows that it is not *simply* a matter of causal relations. A lack of causal relations would not explain why, if we care about comparison, we do not care or, care less, about comparisons between certain times.¹⁵ The relationship between the Earthquake case and the temporal case merits further consideration, but must wait until another time. I note only that differing claims to the resources provides a way to distinguish these cases that causation does not.

¹³ This is not to suggest that there are not also other reasons — like the potential to increase total well-being in the long-run — to develop these capacities.

¹⁴ Chapter 3.

¹⁵ This does not preclude the possibility, however, that it is a further bad for an inequality to be *caused* by one party (i.e. cases where the better-off are responsible for the existence of the inequality).

FORMULATING WEIGHTED PROGRESSIVE EGALITARIANISM

Weighted Progressive Egalitarianism (WPE) results from two principles:

Progressive Scope Restriction (PSR) An individual cannot have a complaint against those individuals who exist entirely at times after which he exists (i.e. who have no temporal overlap with him *and* who live after the time at which he exists).

Temporal Comparison Thesis (TCT) An individual's complaints against those who are temporally distant will have less weight than his complaints against those who are temporally close or contemporaneous.

In Part I, I argued that a theory that satisfied these principles would provide the appropriate rankings of the temporal inequalities. Part II went some length in developing a theory of claims that, if correct, provides justification for this view, as argued in Chapter 9.

These two principles alone, however, still leave much undecided about the exact way to measure complaints. The purpose of this chapter is to give shape to WPE, and evaluate the most plausible options. I will not discuss the problem of how various measures, compatible with the arguments provided in 9, might be justified beyond intuitions; that is work for another day. The goal is to become clear about certain requirements that a plausible theory of WPE must meet.

10.1 COMPLAINTS

Recall that there were several questions which needed to be answered to determine the form of the complaint:¹

1. *Who* has a complaint?
2. What are their complaints *relative to*?
3. *How* are these complaints to be weighted?

We have now made some progress in determining some of the answers.

First, we have learned that in order for someone to have a complaint, they must have a claim in respect of which they have been treated unfairly (comparatively speaking).² Is it also the case that anyone who has been so treated has a complaint?

It is difficult to see why they would not. This is not to say that their complaint would be large — we have not yet determined the form of the complaint — but it seems hard to see that would not have one. Being treated unfairly has been taken to be a personal bad. Therefore I will assume the following principle:

¹*Supra* 34.

² In Chapter 8, I observed that there may be a *noncomparative* aspect to fairness. I set this aside for the remainder of the thesis. Here, when I use the term “unfair”, it is referring to a *comparative* unfairness.

Complaint Ownership An individual has a complaint *iff* he has a claim in respect of which he has been treated unfairly.

Second, we have discovered two principles which restrict the answer to the second question about what complaints are relative to. The first principle is Complaints and Claims (from 8.2). According to it, an individual may only have a complaint against those who have that to which the individual has a claim in respect of which the individual has been treated unfairly. The second principle is PSR. This principle was derived from Complaints and Claims. According to it, an individual may only have complaints relative to his contemporaries or those who live at earlier times.

For purposes of the remainder of our discussion, I am going to continue with the assumptions about the available resources and how they contribute to well-being presented in 9.1.³ As a result, the amount of well-being an individual has is proportional to the amount of resources that individual has. I will also continue to assume that individuals are both equally noncomparatively deserving of D_A , and equally comparatively deserving (of well-being). As a result, comparing their well-being gives some indication of how they have been treated with respect to their claims — namely, by providing an indication of how many resources they have received. In this way, we are back to comparing distributions of well-being, since these reflect the underlying satisfaction of claims.

Given these assumptions, an individual's complaint will be relative to those who are better-off, subject to the condition of PSR. There remains an open question, however, about which particular comparison determines the complaint base. When someone has been treated unfairly with respect to a claim, we know that their complaint base is the difference between their well-being and that of the individual who received that to which they had a claim. But when their claim is disjunctive⁴, who is this person? Or is it many people? That is, do we measure the complaint

³*Supra* 188.

⁴*Supra* 151.

relative to the *best-off person* (BOP) who has that to which they have an unfairly treated claim? Or do we measure the complaint relative to *all those better off* (ATBO)?⁵

It is difficult to decide between these views. On the one hand, it is not necessarily the best-off person who has that which an individual had a claim, since there may be many individuals who are better off. Similarly, it is not necessarily all those who are better off who have that to which an individual has a claim. Someone who is worse off might have some of that to which the individual has a claim, even though he has not been treated unfairly with respect to the worse-off individual. The reason this is so in the case at hand is that his derivative claim to the resources is disjunctive: his claim is denied just in case he does not receive any of the objects in the disjunction. Given our assumptions about the resources, someone who is worse-off but not without will have at least part of the materials that form one of these disjunctions. So unfortunately our theory of claims still leaves much about the complaint base unanswered. Still, we might think that he has not been treated *unfairly* with respect to those who are worse off, so we may safely assume that his complaint is not against them. The answer then is BOP, ATBO, or some hybrid. I'll set the latter possibility aside to focus on just BOP and ATBO. But how are we to decide between them?

Fortunately, what we have learned about the answer to the third question may shed some light on this problem. This is the TCT: the thesis that an individual's complaints against those who are temporally distant have less weight than his complaints against those who are temporally close. Modifying the static versions of BOP and ATBO to satisfy the TCT reveal BOP to be an unsatisfactory position. It also yields a number of lessons for developing the theory of WPE beyond this thesis.

⁵ Alternatively, one might measure the complaint relative to the average, all those above the average, a threshold of some sort, or all those above some threshold. I set these aside.

10.2 THE BEST-OFF PERSON (BOP) & THE LESSON OF G-TYPE HISTORIES

As noted earlier, BOP takes the complaint base to be the difference between the well-being of the best-off person and that of the individual with the complaint. In the static case, this is formalized:

$$c_{i,j} = c(\max(w) - w_{i,j}) \quad (10.1)$$

where $w_{i,j}$ is the well-being of $p_{i,j}$ and w is a vector (a list) of the well-being of the individuals in the distribution ($w = \langle w_{1,1}, \dots, w_{k,n} \rangle$).

Given PSR, this will be further restricted. Rather than being a function of w , $c_{i,j}$ will be a function of the well-being of the members of $p_{i,j}$'s generation g_i , and the well-being of the members of all the generations prior to g_i . We can express this as:

$$c_{i,j} = c(\max(w_1, \dots, w_i) - w_{i,j}) \quad (10.2)$$

where w_i is the vector of the well-beings of the members of generation g_i ($w_i = \langle w_{1,1}, \dots, w_{i,n} \rangle$).

To accommodate the TCT, we need to further modify BOP. An obvious way to do this is to weight the original complaint for time. To do so, we'll need some sort of discount function that takes as its inputs the temporal location of each individual in question. For example:

$$f(t_i, t_j) = e^{\frac{-|t_i - t_j|}{\tau}} \quad (10.3)$$

where t_i and t_j are the temporal locations of individuals p_i and p_j respectively, and τ is a time constant to be determined. This function is equal to 1 when both individuals exist at the same time ($t_i = t_j$) and approaches 0 as the temporal distance between the individuals increases (i.e. as $|t_i - t_j|$ increases). Adjusting

the time constant, τ , affects the rate of this drop-off (i.e. affects the discount rate). There obviously remains much to be worked out regarding the appropriate discount function, but this one will serve our present purposes.⁶ The resulting complaint will then take the form:

$$c_{i,j} = f(t_i, t_{\text{BOP}})c(w_{\text{BOP}} - w_i) \quad (10.4)$$

where t_{BOP} and w_{BOP} are the time and well-being, respectively, of the best-off person p_{BOP} , and where p_{BOP} is the individual whose well-being is given by:

$$w_{\text{BOP}} = \max(w_1, \dots, w_i) \quad (10.5)$$

This modification yields the correct result for the case shown in Figure 10.1, History F. Even though there is the same difference in well-being between $p_{1,4}$ (purple) and the best-off person $p_{1,1}$ (green) as there is between $p_{20,4}$ (blue) and $p_{1,1}$, the complaints are not the same: $p_{1,4}$ is temporally closer to $p_{1,1}$ than $p_{20,4}$ is, and so $p_{1,4}$'s complaint is bigger (matters more).

But while our modification yields the correct judgment in History F, it remains an inadequate fix in the face of more complicated distributions. Consider History G, shown in Figure 10.2. In History G, there are now three more individuals who are better-off with well-being 95 (individuals $p_{20,1}$, $p_{20,2}$, and $p_{20,3}$, shown in lighter green). However, they are not the best-off; that distinction remains with $p_{1,1}$ (darker green).

How does this change affect complaints? According to our modified version of BOP, not at all, with the exception of those now better-off. This is problematic for two reasons:

⁶ For example, is this the appropriate way of discounting? And if it is, what should τ be?

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	100	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	10	10	10	10
g_9	10	10	10	10
g_{10}	10	10	10	10
g_{11}	10	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	10	10	10	10
g_{18}	10	10	10	10
g_{19}	10	10	10	10
g_{20}	10	10	10	10

Figure 10.1: History F

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	100	10	10	10
g_2	10	10	10	10
g_3	10	10	10	10
g_4	10	10	10	10
g_5	10	10	10	10
g_6	10	10	10	10
g_7	10	10	10	10
g_8	10	10	10	10
g_9	10	10	10	10
g_{10}	10	10	10	10
g_{11}	10	10	10	10
g_{12}	10	10	10	10
g_{13}	10	10	10	10
g_{14}	10	10	10	10
g_{15}	10	10	10	10
g_{16}	10	10	10	10
g_{17}	10	10	10	10
g_{18}	10	10	10	10
g_{19}	10	10	10	10
g_{20}	95	95	95	10

Figure 10.2: History G

First, even though $p_{20,4}$ now finds himself the worst-off member of his *own* generation, his complaint remains unchanged from his complaint in History F. This is so even though he is in almost the same situation as $p_{1,4}$ (since the gap is nearly the same size) and arguably in a worse situation (since he is the only one in his generation so greatly disadvantaged).

Second, the weight given to the complaint of $p_{20,4}$ is less than the weight given to the complaint of $p_{10,4}$ (yellow), even though $p_{10,4}$ is located more than 500 years from anyone who is better-off. How large this difference in weighting is depends on the temporal weighting function (and, in our example, the value of time constant τ), but there will still be this difference as long as the weight given continues to decrease over increasing temporal distances.

These two intuitive problems reflect a deeper one: the underlying claims. $p_{20,4}$'s derivative claims to resources during the time at which he lives are the strongest, since these are the only derivative claims to resources which are not backtracked. By comparison, his claim to the resources at the time of $p_{1,1}$ are far weaker. Furthermore, given that individuals cannot have derivative claims to resources located after the time at which they exist, the only individuals who have claims to the resources at the time of $p_{20,4}$ are his contemporaries (as depicted in this distribution). This makes his complaint against his contemporaries strongest: there is no one else who has been treated so unfairly with respect to those derivative claims.

In fact, these problems occur not just in the case of History G above, but in the case of *all* histories like History G, called *G-Type Histories*, with these three features:

- The best-off person is in the first generation.
- There are some *very* well-off people who have just shy of the best-off person located on the opposite end of history.

- Everyone else, in both the first and last generation and all those in between, are much worse-off than any of the better-off and share the same comparatively low well-being.

The problem occurs in these cases because the temporal distance relations between $p_{20,4}$, $p_{10,4}$, and the best-off person, $p_{1,1}$ —the features that determine the complaint on the modified BOP view—remain unchanged regardless of other changes in the distribution (as long as they do not change where the best-off person is). While this version of modified BOP assigns the same complaint to all the worst-off members (the third group), it does not accord them the same weight. Supposing that the gap between the very well-off and the best-off is small, a slight increase for the well-off or a slight decrease for the best-off could reverse entirely whether the complaint of $p_{20,4}$ is the weakest in the distribution, or the strongest. Moreover, in such histories, it will always be the case on the modified BOP view under consideration that more weight will be given to the complaint of an individual in a middle generation distant from either group of the better-off than will be given to the complaint of the worse off in the last generation with the very well-off. Indeed, with our current method of discounting, if the distance is great enough, *negligible* weight will be given to the complaints of those on the very well-off end of history. If temporal distance matters, this seems an absurd way of taking it into account. The existence of some exceptionally well-off individual at a time long past should not influence the size of an individual's complaint, particularly when he also suffers unfairness relative to his contemporaries.

We must abandon this method of modifying BOP. But this exercise has not been without value. Understanding why this modification fails reveals an important lesson for TCT-type modifications. The problem arose because the features on this view that determine the complaint base are not temporally localized. That is, under

the above modified BOP, everyone’s complaint is measured against the well-being of *one* person who is not necessarily temporally local. Thus, problems arise in histories where there are individuals who are at a great temporal distance from the history’s best-off person but who suffer from local inequalities that are almost as severe: their complaints were still only measured against the temporally distant individual — and then discounted for that distance. Another way of looking at it is that while it would be correct, under the TCT, to say that the particular complaint of $p_{20,4}$ against $p_{1,1}$ matters less than the complaint of $p_{1,4}$ against $p_{1,1}$, to use this as the sole determinant of $p_{20,4}$ ’s overall complaint base is to ignore that $p_{20,4}$ may have complaints against those better-off who are local even if they are not *as* well off as $p_{1,1}$. This is the Lesson of G-Type Histories.

This lesson leaves us with two options. The first is to find some way of “localizing” the best-off person against whom complaints are measured. The second is to move away from BOP towards a view that measures complaints against *all-those-better-off*. I will explore each in what follows.

10.3 “LOCALIZING” THE BEST-OFF PERSON (BOP)

If BOP is to be a viable way to measure complaints on WPE, we must find a way to select a temporally “local” best-off person.

The most obvious — and perhaps simplest — way to do so is to use another scope restriction, that is, to find a way to delineate the boundaries of what is temporally “local”. Call the times that are temporally local to an individual his *temporal locale*. An individual’s complaint would then be measured against the best-off person in his temporal locale, that is, the best-off person who is temporally local.

How might such boundaries be determined? One possibility is to determine boundaries by temporal overlap: an individual’s *temporal locale* is restricted to only those who are alive during at least some of the same times as he. With our sim-

plifying assumption of non-overlapping generations, this would be an individual’s generation. This might seem too restrictive: as discussed earlier, it seems plausible that complaints could be held against those in at least a few of the previous generations, since it seems plausible that derivative claims should be able to backtrack at least somewhat. If this is so, then another option might be sought on which an individual’s temporal locale extends beyond the temporal boundaries of his own life towards the past. Such an option would not define an individual’s temporal locale by temporal overlap, but by some other criteria. The question, of course, is what? In particular, how many years or generations should be included if overlap is not the criterion?

I set these problems aside. Suppose that we have found a plausible way for determining the boundaries of a temporal locale. Two questions emerge: First, should complaints still be weighted (at least somewhat) for the distance between the individual and the best-off person within their locale? Second, what are we to do about boundary cases?

The first question is not as benign as it might appear at first glance. If temporal locales are extended, then problems like that presented by History F emerge — the kind of history that pushes in the direction of weighting individual claims for temporal distance. But weighting in the manner discussed above — by distance from the best-off person within the temporal locale — falls into the trap set by G-Type Histories.

This might seem to push in the direction of having smaller temporal locales. But this would make the problem of boundaries more acute. Which raises the second question: what are we to do in boundary cases? Consider History H-1 (Figure 10.3). Suppose that temporal overlap *is* the way to delineate an individual’s temporal locale. Then individual $p_{9,4}$ (blue), and indeed none of his contemporaries (light green), will have any complaint regarding inequality. This is because they are all

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	100	100	100	100
g_2	100	100	100	100
g_3	100	100	100	100
g_4	100	100	100	100
g_5	100	100	100	100
g_6	100	100	100	100
g_7	100	100	100	100
g_8	100	100	100	10
g_9	10	10	10	10
g_{10}	100	100	100	100
g_{11}	100	100	100	100

Figure 10.3: History H-1

equally as well off as any other member of their locale, even though nearly everyone else in their history is *much* better-off. But while it might seem reasonable for $p_{8,4}$ (purple) to have a stronger complaint than $p_{9,4}$, is it reasonable that $p_{9,4}$ should have no complaint at all? Most of the generation immediately preceding $p_{9,4}$ is much better-off. Even though there is no temporal overlap between any of the individuals, they are close enough in time that there easily could have been if $p_{9,4}$ had been born any earlier or if any members of g_8 had died any later. *Prima facie*, they seem to be relevantly temporally close. As noted earlier, they seem at least relevantly close for the backtracking of claims.⁷ If this is correct, then a temporal locale restricted to one generation or mere temporal overlap is too restrictive.

Suppose, then, that we expand the notion of temporal locale to two generations — one’s contemporaries, and the generation immediately preceding one’s own. Then $p_{9,4}$ would have a complaint — one the same size as $p_{8,4}$ — measured against the best-off person in this temporal locale, which would be one of the first three members of g_8 . We can set aside the question of whether to further weight for distance in this case, for it seems arguments could be made in either direction as long as the

⁷*Supra* 195.

discounting were not extreme: this size temporal locale does not seem to be large enough to run into a History G type problem. So far, so good.

Now consider History H-2 (Figure 10.4). History H-2 is similar to H-1. g_8 and

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	250	250	250	250
g_2	250	250	250	250
g_3	250	250	250	250
g_4	250	250	250	250
g_5	250	250	250	250
g_6	250	250	250	250
g_7	250	250	250	250
g_8	100	100	100	10
g_9	10	10	10	10
g_{10}	100	100	100	100
g_{11}	100	100	100	100

Figure 10.4: History H-2

g_9 enjoy the same levels of well-being in each. The difference is that all members of g_7 and earlier are now even better-off than they were in H-1, enjoying well-being of 250. This affects the complaint of $p_{8,4}$, for the best-off member of his temporal locale is now one of the individuals in g_7 . The complaint of $p_{9,4}$ remains the same, measured against one of the better-off individuals of g_8 .

That the complaint of $p_{9,4}$ remains unaffected might strike some as counterintuitive. Those best-off members of his world are just outside the boundaries of his temporal locale, a mere one generation between them. Perhaps, one might think, his complaint is not as *strong* as someone who is closer, but surely he *does* have a complaint regarding inequality against them.

How are we to decide? First, it might be useful to consider the complaints of the other members of $p_{9,4}$'s temporal locale, and, in particular, the complaints of those better-off members of g_8 . The temporal locale of individuals in g_8 includes the members of g_7 . The best-off members of $p_{9,4}$'s temporal locale are among the

less well-off within their own temporal locale. Indeed, $p_{8,1}$, $p_{8,2}$, and $p_{8,3}$ each have a complaint base of 150. This is larger than $p_{9,4}$ ’s complaint base of 90, implying that $p_{8,1}$, $p_{8,2}$, and $p_{8,3}$ suffer a greater unfairness or slight regarding inequality. Is this plausible?

Perhaps, the driving intuition behind this judgment is not that $p_{9,4}$ has a complaint against the members of g_7 but that his complaint matters more than those of $p_{8,1}$, $p_{8,2}$, and $p_{8,3}$ because he is worse off than they are. Weighting for levels of well-being might help address this concern, then, to a point. But we could alter the case again in such away as to increase the well-being of g_7 to arrive at a point where $p_{9,4}$ ’s complaint is again outweighed by that of his best-off persons $p_{8,1}$, $p_{8,2}$, and $p_{8,3}$. Perhaps, at this point, the difference would indeed be plausible because of how much worse-off $p_{8,1}$, $p_{8,2}$, and $p_{8,3}$ are relative to g_7 . But it does seem bizarre that even in this case, $p_{9,4}$ ’s complaint remains *completely unaffected*. If one was already indecisive about whether $p_{9,4}$ should have at least some complaint regarding inequality against the members of g_7 , this comparison with the complaints of the better-off members of g_8 might pack some punch.

A better way to draw this out might be to consider a case more similar to H-1. Call this H-3, depicted in Figure 10.5. History H-3 is like History H-2, but the best-off people in $p_{9,4}$ ’s temporal locale have been lowered to the same level of well-being as $p_{9,4}$. Thus, there is perfect equality within $p_{9,4}$ ’s temporal locale, and so, according to this temporal locale version of BOP, $p_{9,4}$ has no complaint regarding inequality. This is so even though he is in one of the two worst of generations in the history as shown, and even though a generation only one removed – g_7 — enjoys a level of well-being many multiples greater than his own. While it is plausible that $p_{9,4}$ would have a stronger complaint if those better-off people were temporally closer, that he should have *no* complaint in this case seems absurd. There would appear to be an abundance of resources within the scope of his backtracking derivative claims,

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	250	250	250	250
g_2	250	250	250	250
g_3	250	250	250	250
g_4	250	250	250	250
g_5	250	250	250	250
g_6	250	250	250	250
g_7	250	250	250	250
g_8	10	10	10	10
g_9	10	10	10	10
g_{10}	100	100	100	100
g_{11}	100	100	100	100

Figure 10.5: History H-3

and, if so, he is being treated unfairly with respect to those claims.

This might lead us to once again expand the temporal locale. But this will lead to yet another iteration of the problem of H-Type Histories. Expanding the locale to be large enough to avoid this problem will then reintroduce the problem of History F, and any attempt at weighting within the temporal locale for distance will lead to the problem created by G-Type Histories. What to do?

Another way of trying to approach the problem might reveal a lesson. Consider History I (Figure 10.6). Who is the best-off person against whom $p_{9,4}$'s complaint should be measured? The way that History I is structured, whoever one chooses as the best-off, they will have a *much* greater complaint against someone within their own temporal locale. Continuously expanding the boundaries of $p_{9,4}$'s temporal locale will not alleviate this problem. And it will create the pressure, mentioned before, of falling into the problem of weighting posed by History F and the G-Type Histories.

The temporal locale approach creates great problems by drawing such sharp boundaries. So this method of revising BOP does not appear promising since the sharp scope restrictions of temporal locales called for do not appear promising.

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	2000	100	100	100
g_2	1500	100	100	100
g_3	1100	100	100	100
g_4	800	100	100	100
g_5	550	100	100	100
g_6	350	100	100	100
g_7	200	100	100	100
g_8	100	10	10	10
g_9	10	10	10	10
g_{10}	100	100	100	100
g_{11}	100	100	100	100

Figure 10.6: History I

This might appear to be a reason to abandon BOP, or the intuitions underlying TCT. Fortunately, this is not yet necessary. It seems plausible that our method for determining the size of $p_{9,4}$'s complaint on BOP needs to be more nuanced than the delineation of a crisp boundary of potential candidates for the relevant best-off person. I now turn to one such possibility.

10.4 A DIFFERENT WAY TO MODIFY BOP: TEMPORALLY SELECTIVE BOP

An alternative way to determine the relevant best-off person is to determine which one would provide the greatest complaint given the temporal weighting. This complaint might be seen as the dominating one, and is determined by the combination of the temporal distance and the absolute size of the welfare gap. Call this view *Temporally Selective BOP*. The complaint of person $p_{a,b}$ might then be formalized:

$$c_{a,b} = c(\max(f(t_{a,b}, t_{i,j})(w_{i,j} - w_{a,b}))) \quad \text{for } i \leq a \quad (10.6)$$

where a is between 1 and k (for k generations) and b is between 1 and n .

This way of modifying BOP avoids the problems encountered earlier. First, it weights complaints in a way that yields the plausible result found earlier with respect to History F. In that case, where there was only one best-off person, the complaints of individuals will be weighted according to their distance from that individual.

Second, this view respects the Lesson of the G-Type Histories by allowing local inequalities to affect an individual's complaint even if they are temporally distant from the best-off person in their history. Recall History G (Figure 10.2). According to *Temporally Selective BOP*, the complaint of $p_{20,4}$ is now dominated by the local inequality. Since the best-off member of $p_{20,4}$'s generation is almost as well-off as the best-off member of $p_{1,1}$'s generation, $p_{20,4}$'s complaint is *almost* as big as that of $p_{1,1}$. Moreover, the complaint of $p_{10,4}$ will not only be less than that of $p_{20,4}$ (a problem in the previous G-Type History) but it will be much less. This result fits well with the initial intuitions about this case.

Third, all this has been achieved without resorting to the sharp scope restrictions used in the last section. In such a way, the problems encountered in the H-Type Histories might be avoided. Boundary cases are not determined by mere temporal distance, but by a combination of the temporal distance and the welfare gap. Whether this would completely alleviate the problems encountered above depends on the form of $f(t_i, t_j)$, but this is a question we must leave for another time.

This leaves a final question about whether this modification is in keeping with the spirit of BOP. It is clear that it is not a direct version of the BOP view: Temporally Selective BOP cannot be arrived at by identifying the best-off person and then measuring the complaint base relative to that person. Rather, the formula is arrived at by assessing the difference relative to each better-off person within the distribution and then selecting the greatest one after discounting has been applied. In this way, it might be called a version of ATBO, the view that takes the complaint to be

measured relative to *all-those-better-off*.⁸

But in another way, it is in keeping with the spirit of BOP: the process of determining the complaint may include comparing differences relative to all-those-better-off, but the complaint base itself is taken relative *only* to the largest of these once discounted for time. In this way, Temporally Modified BOP looks at all the comparisons, but selects the greatest offender to act as the relevant “Best-off Person”. This is very different from versions of the ATBO view which take the complaint base as some sort of aggregate of *all* such differences between the individual and those better-off (e.g. a summation of pair-wise complaints).

Earlier, I noted that the Lesson of the G-Type Histories left us with two options: find a way to “localize” the best-off person against whom complaints are measured, or turn to a measure that takes complaints as relative to *all those better-off*. Having now discussed (and exhausted⁹) the ways of modifying BOP that so “localize” the best-off person and accommodate the TCT, it is time to consider the second option.

10.5 TEMPORALLY MODIFIED ALL THOSE BETTER-OFF (ATBO)

The basic thought behind the relative to *All Those Better-Off* (ATBO) view of complaint measurement is this: if it is bad for there to be one person who is better-off than me, then it is even worse if there are many people who are better-off than me.¹⁰ On ATBO complaints are measured relative to all those better-off than an individual. There are many ways in which this might be done. The most straightforward is to simply sum the differences between the well-being of the individual and each of those

⁸ John Broome called my attention to this.

⁹ Perhaps there are others, but I cannot think of any.

¹⁰ Temkin also makes this observation in presenting ATBO. (Temkin (1993, 26))

who are better-off. Combined with PSR, the complaint for person $p_{a,b}$ is given:

$$c_{a,b} = \sum_{i=1}^a \sum_{j=1}^n \max(0, w_{i,j} - w_{a,b}) \quad (10.7)$$

where a is between 1 and k (for k generations) and b is between 1 and n .

This simple version is also straightforward to modify for compliance with the TCT. Since comparative complaints matter less the greater the temporal distance between the individuals in question, the contribution of each of these differences in well-being would be weighted by our discount function. Call this *Temporally Modified ATBO*.

$$c_{a,b} = \sum_{i=1}^a \sum_{j=1}^n f(t_{a,b}, t_{i,j}) (\max(0, w_{i,j} - w_{a,b})) \quad (10.8)$$

This way of measuring complaints and weighting them for temporal distance yields plausible results for the problematic cases outlined above.

In History F, the complaint of $p_{1,4}$ (purple) is greater than that of $p_{20,4}$ (blue) because of the temporal distance.

In the G-Type Histories (defined on page 212), local inequalities factor very heavily into an individual's complaint, since every pair-wise inequality is accounted for.

In the H-Type Histories, ATBO avoids some of the same problems avoided by the Temporally Modified BOP, including those of the boundary cases, since it does not make use of sharp scope restrictions. Whether History I causes a problem depends on the temporal weighting function chosen, but since the discount is meant to reflect the weakening of backtracking claims, it seems plausible.

10.5.1 POTENTIAL PROBLEMS WITH TEMPORALLY MODIFIED ATBO

Might there be distributions that are problematic for Temporally Modified ATBO? The best place to begin is to consider those cases that are problematic for the unmodified ATBO, and then consider whether these problems would be exacerbated — or alleviated — by weighting for time.

Consider the static case, depicted in Figure 10.7, that some find problematic on ATBO. Suppose that each of the two distributions depicted have the same number



Figure 10.7: Static Cases J & K

of people (80, the same number as the other cases considered). In J, the best-off individual has a well-being of 100 and the rest have well-being 50. In K, the best-off have well-being 70 and the worst-off individual has well-being 65. Who should have the larger complaint regarding inequality — one of the worst-off in J, or the worst-off member of K?¹¹

On one view discussed thus far, BOP, the complaint of one of the worst-off individuals in J would be greater since the complaint base is the size of the gap between an individual and the best-off person. According to BOP, the complaint base of a worse-off member of J is 50, while the complaint base of the worst-off member of K is 5. But on ATBO, the worst-off member in K has a greater complaint because there are so many people better-off than he even though they are not better-

¹¹ This is a different question from this one: *which* inequality is the worse inequality, J or K? ATBO, when combined with an Additive Principle of Equality (AP), agrees with BOP&AP that J is the worse *distribution*.

off by that much: his complaint base would be 395, as contrasted with the complaint base of a worse-off individual in J of 50. Some might worry that this is implausible. But is it?

A critic of ATBO could argue that the problem is the straight summing of pair-wise differences to arrive at the complaint base. This type of summation depends on the assumption that each additional person who is better-off than an individual by a given amount generates the same amount of disvalue. A critic might deny this assumption, arguing that there might be diminishing marginal disvalue for each better-off person. That is, while it is worse for there to be two people better-off than an individual by a given amount rather than only one person better-off than that individual, it is not twice as bad. This would map well the theory of claims: the individual had a singular claim which was denied by being denied a number of particular resources, but not by being denied all those resources.

Alternatively, a critic could argue that the problem lies in the relation between the worse-off groups, the better-off groups, and the average. While the worst-off in K is farther below the average of his distribution than the worse-off members of J are below their average, the difference is small. But the better-off in K are barely above the average while the best-off in J has nearly twice the average.

The solution in the static case is reasonably simple if the concern is the first just presented: the complaint base will not be a straight summing of the pair-wise differences, but will allow for a decreasing contribution of each comparison of the individual's well-being to that of a person who is better-off. If the concern is the second, the solution in the static case is less clear if an ATBO view is to be maintained.

I set these aside. Whether this is indeed a problem and which solution (if either) is appropriate are questions for the static egalitarian debate. Supposing they might

be resolved, our concern is whether something similarly problematic might appear in the extension of temporally modified ATBO over time.

Consider the following two histories, History J (Figure 10.8) and History K (Figure 10.9), modelled after the static case. In History J, there are 80 people

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	100	50	50	50
g_2	50	50	50	50
g_3	50	50	50	50
g_4	50	50	50	50
g_5	50	50	50	50
g_6	50	50	50	50
g_7	50	50	50	50
g_8	50	50	50	50
g_9	50	50	50	50
g_{10}	50	50	50	50
g_{11}	50	50	50	50
g_{12}	50	50	50	50
g_{13}	50	50	50	50
g_{14}	50	50	50	50
g_{15}	50	50	50	50
g_{16}	50	50	50	50
g_{17}	50	50	50	50
g_{18}	50	50	50	50
g_{19}	50	50	50	50
g_{20}	50	50	50	50

Figure 10.8: History J

spread evenly over 20 generations. The best-off person ($p_{1,1}$, in green) has well-being level 100, and lives during the first generation. Everyone else has well-being 50.

In History K, there are likewise 80 people spread evenly over 20 generations. The worst-off person has well-being 65, and lives during the last generation. Everyone else has well-being 70.

The only difference between these histories and the related static case are the location of the individuals in time. Do similar problems arise?

	$p_{i,1}$	$p_{i,2}$	$p_{i,3}$	$p_{i,4}$
g_1	70	70	70	70
g_2	70	70	70	70
g_3	70	70	70	70
g_4	70	70	70	70
g_5	70	70	70	70
g_6	70	70	70	70
g_7	70	70	70	70
g_8	70	70	70	70
g_9	70	70	70	70
g_{10}	70	70	70	70
g_{11}	70	70	70	70
g_{12}	70	70	70	70
g_{13}	70	70	70	70
g_{14}	70	70	70	70
g_{15}	70	70	70	70
g_{16}	70	70	70	70
g_{17}	70	70	70	70
g_{18}	70	70	70	70
g_{19}	70	70	70	70
g_{20}	70	70	70	65

Figure 10.9: History K

On Temporally Modified ATBO, the complaint of the worse-off in History K is less than it would be in static distribution K. This is because the comparisons that form her complaint base vary according to the temporal distance. The result is that the comparisons against the better-off do not all generate the same disvalue. Those that are temporally local have greater disvalue, while those that are temporally extended count for very little. Exactly what the size of her complaint would be depends on the chosen discount, but it is clear that it would be smaller than in distribution K, without a need for revision of the static ATBO theory.

How the complaint of the worse-off in History K compares to that of the worse-off in History J varies by worse-off individual. It will be closer in size to the complaint of $p_{1,4}$ (purple), but likely larger than that of $p_{10,4}$ (yellow) and $p_{20,4}$ (blue). This seems a more palatable result than what ATBO yielded in the static case. Temporally Modified ATBO appears to escape the original problem in the static case, at least for fixed populations extended over time.¹²

10.6 CONCLUSIONS: LESSONS LEARNED MOVING FORWARD

Three lessons emerged in our discussion:

1. The Lesson of the G-Type Histories
2. The Lesson of the H-Type Histories
3. The Lesson of the I-Type Histories

¹² Weighting for time seems to solve at least one of the problems ATBO would encounter in cross-temporal comparatives when extending the theory to cover different number distributions. You could keep adding more and more better and better-off people at earlier times without complaints growing too tremendously — or as tremendously — as if they were all added to the same population-at-a-time.

Moreover, given PSR, you could keep adding more and more better and better-off people at later times without influencing the complaints of earlier individuals *at all*.

The Lesson of the G-Type Histories is that the feature determining the complaint base must be temporally localized.

The lesson of the H-Type Histories is that sharp scope restrictions to certain regions of time or “temporal locales”, thereby excluding rather than discounting comparisons to areas outside the locale, are in conflict with the nature of backtracking claims. Attempts to modify the boundaries of the temporal locale to accommodate seemingly relevant comparisons reintroduce the problems of the F-Type Histories. But attempts to alleviate the F-Type problem using discounting within the locale only led to a reiteration of the problem created by G-Type Histories. A plausible theory of Temporally Weighted Egalitarianism must avoid sharp scope restrictions to temporal locales beyond PSR.

We ended with two plausible ways to formalize WPE that satisfy these problems: Temporally Selective BOP and Temporally Modified ATBO. There remained a question, however, of whether they would meet the challenge of the I-Type Histories. How these views answer this case depend in part on the discounting function. Further exploration is required, but is beyond the scope of this work.

WEIGHTED PROGRESSIVE EGALITARIANISM: IMPLICATIONS

This final chapter explores two *prima facie* advantages of WPE that provide reason to develop the theory further. First, WPE provides a rebuttal to the argument from overlapping generations for TNE. Second, Weighted PE seems poised to help alleviate one problem in population axiology. We shall discuss each in turn.

11.1 THE ARGUMENT FROM OVERLAPPING GENERATIONS

Our discussion, for the most part, has proceeded on the assumption of non-overlapping generations. That is to say, individuals either had complete temporal overlap or none at all. In so doing, we avoided several difficult questions about equality and time, in particular the question about the appropriate temporal unit for egalitarian concern (e.g. whole lives, simultaneous segments, or corresponding segments) and the

argument for TNE from overlapping generations.¹

The argument for TNE from overlapping generations runs as follows. Consider a case with 3 individuals from partially overlapping generations, pictured in Figure 11.1. Person A has partial overlap with person B and person B has partial overlap with person C, but person A and person C have no temporal overlap.




Figure 11.1: Overlapping Generations

Suppose that equality matters between overlapping individuals, even if the overlap is only partial. Make the further assumption that the proper object of egalitarian concern is equality of individuals' lifetime well-being.² Given these two assumptions, equality between A and B matters, and equality between B and C matters. Since equality is a transitive relation, the argument goes, equality between A and C matters.

In other words, it is bad for A to be worse off than B, and for B to be worse off than A, through no fault of their own. Likewise, it is bad for B to be worse off than C, and for C to be worse off than B, through no fault of their own. Since this entails

¹ Whole lives, simultaneous segments, and corresponding segments are the three options proposed by McKerlie in McKerlie (1989). The problem of equality between individuals with partial temporal overlap has long been a subject of debate in discussions of intergenerational justice. While weighing inequality in such cases is important, a detailed examination of cases of partial overlap requires more space than we have here available. See also Laslett and Fishkin (1992)

² The view I have defended is more complicated than this — it is concerned about the equal satisfaction of claims, but measures the unfairness of inequality by comparing well-being. Thus far, I have assumed that the appropriate measure is lifetime well-being, but I have not argued for this position.

that it is good for $w_A = w_B$ and for $w_B = w_C$, then, by transitivity, it is good for $w_A = w_C$. So, if equality of lifetime well-being is what matters, and equality between people with temporal overlap matters, then equality between nonoverlapping people also matters — or so the argument goes.

This argument is something of a strawman. The leap from assuming that equality of lifetime well-being matters between partially overlapping generations to TNE is just that: a leap. The argument seems to confuse the transitivity of certain comparisons (here, equality) with the transitivity of the value of those comparisons. It is true that if equality between A and B, and between B and C, both matter, then the best distribution of a fixed total of well-being will be the one in which all three are equally well-off. But this is not because the comparison between A and C matters; rather, it is a byproduct of minimizing the disvalue of the comparison between A and B while also minimizing any disvalue from the comparison between B and C.

To illustrate, compare three possible outcomes:

1. $w_A = 10$; $w_B = 10$; $w_C = 5$
2. $w_A = 5$; $w_B = 10$; $w_C = 10$
3. $w_A = 10$; $w_B = 5$; $w_C = 10$

(1) and (2) each have one inequality between temporally overlapping individuals. In (1), it is between B and C; in (2), between (A) and (B). In (3), however, there are two inequalities between temporally overlapping individuals: between A and B, and between B and C. If inequalities are measured relative to *all those better off*, then (3) is worse than both (1) and (2). But TNE does not make this distinction, since TNE also counts a second inequality in outcomes (1) and (2), that between A and C.

Very well, the temporally neutral egalitarian might reply, but you have missed the point. The badness of inequality of lifetime well-being is basic: if the comparison between A and B matters, and the one between B and C matters, what is the relevant difference that makes the comparison between A and C matter less or not at all?

This is to change the line of argument, which the temporally neutral egalitarian must, because all this example shows is the very point on which we disagree — and nothing about who is correct. I have already offered my answer to his question above.

Before moving on, it bears noting that the supposed argument above arises from the assumption that a comparison of lifetime well-being is what matters. It is a natural assumption if one assumes that time does not matter. It is also the assumption that leads to TNE — no need for this strawman argument. But, as McKerlie also argued, this is not an obvious assumption. The location of a good in time may not affect whether that good matters, but the times of the relata in a given comparison might affect whether that *comparison* is a good that matters (in which case, it will be a good whenever in time such a relation holds).

Indeed, I suspect McKerlie would not be pleased with WPE's use of the comparison of lifetime well-being as the measure of the disvalue of the unfairness when claims are unequally satisfied. But this should not cause us too much concern. The use of this comparison has only been a simplifying tool, used to try to evaluate other parameters that WPE must meet. The problem is not that WPE takes this as the appropriate measure; the problem is that WPE is not fully developed. There is much work left to be done.

Fortunately, WPE and the theory of claims supporting it seem to provide a foundation for launching into this work about the proper temporal unit of egalitarian concern. WPE provides a reason to take equality between those individuals who are greatly overlapped as at least somewhat more important than inequality between

individuals who have negligible overlap. It could support a variation on a segment view, whereby simultaneous segments are weighted more heavily, but past inequalities can factor in. It could also support a corresponding segment view where equality between corresponding segments of individuals with greater temporal overlap matter more than equality between corresponding segments of individuals with minimal temporal overlap. Such views appear plausible, and are worth further exploration at another time.

11.2 WPE, THE NEUTRAL LEVEL & CONTEXT SENSITIVITY

Throughout, I have assumed a fixed population. Will WPE extend to different number cases? Undoubtedly, navigating it through the paradoxes may be difficult – if not impossible.³ That is not my concern here. Rather, WPE has some interesting implications for extension to different number cases that are worth mentioning. To get to them, however, I need to go by way of a discussion about the neutral level of personal good.

11.2.1 THE NEUTRAL LEVEL

Suppose we have found a satisfactory population axiology – or betterness relation – for these fixed-number populations. One may then ask: for a given population, how will adding a person affect the value of this population?

Many intuitively believe that, for some sizeable range of well-being, adding a person to an existing population will not *itself* affect the value of that distribution. This is the *Intuition of Neutrality*, and it must be rejected. Rather, we must accept the existence of what is called the “Neutral Level”:

³*cf.* Arrhenius (2000).

The Neutral Level (of Personal Good) for a given population, the level of individual good at which adding a life to that population does not affect the value of the population

For a given population, there is only *one* neutral level. This neutral level forms a “sharp boundary” between levels of well-being at which adding a life to a distribution *improves* that distribution, and levels of well-being at which adding a life *worsens* a distribution.⁴

The argument for the neutral level has been forcefully made by Broome⁵; I will not repeat it here. Rather, the conclusion that there is a neutral level leaves two questions: What is the neutral level? And is it independent of the existing population? The first poses a daunting task. I shall instead focus on the second question in the hopes that it might provide a first step towards illuminating the answer to the first.

In the first section, I shall argue, as Broome does, that those who deny the value of comparative complaints must maintain that the neutral level of *individual good* is context-independent – that is, constant for all populations. In the second section, I shall argue that egalitarians (who are pluralists about the good) can maintain that the neutral level of *well-being* depends *directly* on context. In the third section, I raise an objection to the use of TNE for these purposes, one which WPE avoids.

11.2.2 AN ARGUMENT FOR A CONSTANT NEUTRAL LEVEL

INITIAL ASSUMPTIONS

To get any argument regarding the neutral level off the ground, certain assumptions need to be, or have already been, made about the structure of the good. I hope

⁴ Broome (2004, 142)

I use “improving/worsening a distribution” and “increasing/decreasing the value of a distribution” interchangeably.

⁵ Broome (2004, Ch. 14.2, p. 202-208)

that these will not appear too contentious. They are well-defended elsewhere in the literature, and the costs of abandoning them are high.⁶

The first is the Principle of Personal Good⁷ for fixed number populations:

- Principle of Personal Good** 1. Two alternatives are equally good if they are equally good for each person.; AND
2. If one alternative is at least as good as another for everyone, and definitely better for someone, then it is better.

In other words, the Principle of Personal Good is the claim that the value of a fixed population depends solely on the value of individual good, or the good for each individual, and that it is *increasing* in individual good.

The second and third assumptions are related: coherence of the individual betterness relation (or individual good), and coherence of the general betterness relation (or social good). That is, the betterness relations conform to the axioms of expected utility theory.⁸

The fourth assumption is Bernoulli's Hypothesis:

Bernoulli's Hypothesis One alternative is at least as good for a person as another *iff* it gives the person at least as great an expectation of her good.

This assumption is likely to garner the greatest resistance: it is the assumption of risk neutrality about the good. Defending it takes care and will bring us too far afield; I again refer the reader to Broome (1991).

Finally, there is the assumption of impartiality. Formally, impartiality implies that each individual's good is given equal weight.

⁶ For a detailed discussion, see Broome (1991)

⁷ As presented in Broome (1991, 155)

⁸ Many of the objections against such coherence began as objections against the coherence of preference relations, not betterness relations. While such objections are powerful when applied to preferences, they lose their force when applied to betterness relations. See Broome (1991, Ch. 6, 7.3, & 10)

Together, these yield the Same-Number Addition Theorem. First proved by Harsanyi, and later clarified by Broome⁹, the theorem is as follows:

Same-Number Addition Theorem “Of two distributions whose populations have the same size, one is better than the other if and only if it has a greater total of people’s good.”¹⁰

With these assumptions in place, the argument for a constant neutral level of personal good is reasonably straightforward.¹¹

THE ARGUMENT

Broome’s argument that the neutral level of personal good is constant may be outlined as follows:

The first two steps are assumptions already discussed. The third is the assumption that there is a value function that represents the goodness of *every* distribution.¹²

Assumptions

1. “Of two distributions whose populations have the same size, one is better than the other if and only if it has a greater total of personal good. (*The Same Number Addition Theorem*)
2. For a given population, there is a level of personal good, ν , at which it is neither good nor bad if a person is brought into existence at this level. (*Definition: Neutral Level*)

⁹ Broome (1991, Ch. 10.1)

¹⁰ Broome (2004, 137)

¹¹ One could make a similar proof for the constancy of the neutral level of personal good without the assumption of the Bernoulli Hypothesis and resulting Same Number Addition Theorem. Prioritarianism is one such view. On these views, the value of a fixed population is determined by the sum of the transformed good. For simplicity, I set this complication aside.

¹² Perhaps there is not, though I remain hopeful.

3. There is a value function that represents the goodness of every distribution.

The next step is to apply the Same-Number Addition Theorem (*Assumption (3)*) to obtain the value function for populations of two fixed sizes: n , and $n + 1$.

Derivation of Value Functions for Fixed-Size Populations of Sizes n and $n + 1$

4. Applied to distributions of n people, the value function is an increasing transform of the total personal good. (*from the Same-Number Addition Theorem.*)
Notate this $f_n(G_n)$, where G_n is the total wellbeing and $f_n(\cdot)$ is an increasing transform.¹³

5. Applied to distributions of $n + 1$ people, the value function is an increasing transform of the total personal good in distributions with $n + 1$ people. (*from the Same-Number Addition Theorem*) Notate this $f_{n+1}(G_{n+1})$.

Next, compare the value of a given population of n people with the value of another population of $n + 1$ people that is created by adding one person, at the neutral level, to the population of n people:

Comparison of Value of a Given n -size Population with Value of Given Population plus New Member at the Neutral Level

6. Suppose there is a population of size n , with total personal good G_n . The value of this population will be given by $f_n(G_n)$.
7. Suppose a new person is added to the population in (5) at the neutral level, ν .
8. Suppose that this new person does not affect the total good of the original population, G_n .

¹³ An increasing transform preserves an ordering.

9. The new population will be a distribution of $n+1$ people, and will have total good $(G_n + \nu)$.
10. The value of the new population created in (7) is given by $f_{n+1}(G_n + \nu)$. (*from (5)*)
11. The value of the original population with n people, and the value of the population comprised of the original population and the person at the neutral level is the same. That is, $f_n(G_n) = f_{n+1}(G_n + \nu)$. (*from the Definition of the Neutral Level*)

Steps (6) – (11) are the same for any population of size n . From (11), we can see that if the neutral level ν , depends on context, it can only depend in some way on the *total* good, G_n , of the original population.

Broome continues:

What the Neutral Level of Personal Good May Depend On

12. The average level of good of a distribution, A_n , is the total good divided by the number of people, G_n/n . (*Definition: Average Level of Personal Good*)
13. The equation in (11) may be rewritten using the average level of personal good and the number of people in the distribution yielding $f_n(nA_n) = f_{n+1}(nA_n + \nu)$. (*from (11) and (12)*)
14. Given a distribution, the equation in (13) may be solved for ν in that context.
15. If the neutral level of personal good depends on context, it can only depend on the number of people, n , and/or on the average personal good, A_n .

Broome then provides arguments for the following:

The Neutral Level of Personal Good does Not Depend on Context

16. The neutral level of personal good does not depend on the average level of good. (see Broome (2004, Ch. 13.3))
17. The neutral level of personal good does not depend on the number of people. (see Broome (2004, Ch. 13.4))

I will not here rehash his arguments for (16) and (17). The arguments against the average view, and in support of (16), are many, well-known, and convincing.¹⁴ Premise (17) is harder to defend. I will assume it to be correct.¹⁵

From this, he concludes:

18. The neutral level of personal good does not depend on context (i.e. is constant).

If comparative distributions do not matter, then personal good is an individual's well-being (since nothing else contributes to his personal good). In this case, the neutral level of personal good is the level of well-being for which adding a person with that level of well-being neither improves nor worsens a distribution.

WHY THIS IS TROUBLING

The constancy of the neutral level returns us to a very difficult question: what is it?

Placed too low, the neutral level threatens to run population axiology into yet another iteration of the Repugnant Conclusion.¹⁶ The yet-to-exist would come into direct competition with the existent for those resources that contribute to individual good. This will be particularly true if there are diminishing marginal returns of resources to the good and low "entry" costs: maximization of the good might then more strongly favor adding new individuals to the population. A population of high

¹⁴ e.g. Parfit (1984, 420) and Broome (2004, Ch. 13).

¹⁵ If Premise (17) is incorrect, it raises some interesting questions. If the value of a population depends, in a direct way, on the number of people, would it depend on the historic total, or concurrent total?

¹⁶ Originally presented in Parfit (1984)

quality lives would be worse than a much larger population with lower quality lives just at or above the neutral level.

This places pressure to have a higher neutral level. Doing so, however, threatens to run population axiology into a version of the Reverse Repugnant Conclusion: a population of lives full of great suffering would be better than a much larger population with lives just below the neutral level.

These are the two primary problems that make the constancy of the neutral level seem implausible. Broome observes:

The implausibility arises principally from the interaction of the positive and negative repugnant conclusions. To make ourselves comfortable with the positive one, we need a high neutral level, but this makes us uncomfortable with the negative one. We find it hard to believe there is a level of wellbeing such that it is good to add people to the population at any higher level, and bad to add people at any lower level. But this is just the neutrality intuition again: it is implausible that there is just a single neutral level. It is hard to reject this intuition, but we have to. Once we have rejected it, we have to recognize there is a single neutral level. Adding people below this level is a bad thing, and its badness must weigh against other bad things such as suffering.¹⁷

But Broome is wrong on a different point. The point of these objections is separate from the intuition of neutrality; the intuition of neutrality appeared as one way out of them, a way that, as noted, must be rejected. Rather, one might have the intuition that the neutral level depends heavily on context without believing that there is a neutral range. That is, one might have the intuition that the neutral level forms a sharp divide between levels of well-being for which an added life at that level would improve or worsen a distribution while also holding the intuition

¹⁷ Broome (2004, 213)

that the sharp divide is located at different levels depending on features about the population to which the life might be added. Even without the Intuition of Neutrality, the ramifications of a constant neutral level might still be disturbing.

My goal here is neither to solve this problem nor explain the source of the troubling intuitions. Rather, it is to offer these as reason to be dissatisfied with the conclusion that the neutral level of *well-being* is constant. In the end, we may need to accept some counterintuitive or dissatisfying results. But first, it is worth considering whether the neutral level *must* be constant, and if not, whether that would help escape any of these problems. I turn to the first of these questions in the next section.

11.3 COMPARATIVE COMPLAINTS & THE NEUTRAL LEVEL

It will now be useful to make a distinction between two types of Neutral Levels (for existence) alluded to earlier:

The Neutral Level of Personal Good for a given population, the level of personal good at which adding a life to that population does not affect the value of the population.

The Neutral Level of Well-being for a given population, the level of well-being at which adding a life to that population does not affect the value of the population.

For non-comparative views – views in which an individual’s *personal* good is unaffected by the distribution of good or well-being – these will be equivalent.¹⁸

¹⁸ Note: Prioritarianism, or the Priority View, is *not* a comparative view. Personal good may be a concave transform of well-being, but their *neutral* levels are equivalent.

That is, the Neutral Level of Well-being, like the Neutral Level of Personal Good, will be constant.

An individualistic egalitarian must accept that the neutral level of *personal good* is constant, but they need not accept that the neutral level of *well-being* is constant. This is because, on their view, personal good includes a comparative component in addition to an individual's well-being. The neutral level of personal good would likewise include a comparative component.

Recall that according to an individualistic egalitarian, an individual's good is comprised of two components: their well-being (or good excluding the intrinsic disvalue of comparative complaints) and any comparative complaints. This is formalized as:

$$g_i = w_i - c_i \quad (11.1)$$

where g_i is the individual good of person p_i , w_i is their well-being, and c_i is their complaint.

Similarly, the neutral level of personal good is given as a difference between well-being and any comparative complaints:

$$\nu = w_\nu - c_\nu \quad (11.2)$$

where ν is the neutral level of personal good, w_ν is the neutral level of well-being, and c_ν is the value of any comparative complaint the added person has.

While the neutral level of *personal good*, ν , is constant, this does not entail that the neutral level of *well-being*, w_ν , is constant. In fact, the inclusion of a *comparative* complaint will ensure that the neutral level of well-being is dependent on context, with the nature of the dependency determined by the complaint function, c_i .

To demonstrate, consider a simple example. Suppose that all but the best-off have complaints, and that these complaints are measured relative to the well-being of the best-off person:¹⁹

$$c_i = \alpha \max(0, w_{\text{BOP}} - w_i) \quad (11.3)$$

where α is a scaling constant, w_i is the well-being of the individual with the complaint, and w_{BOP} is the well-being of the best-off person.

Consider a population of size n in which everyone shares the same level of well-being, w^* . Call this Population A. Since it is a perfectly equal distribution, there are no comparative complaints. This entails that the personal good of each individual within the distribution, g_i , is equal to their well-being, w^* .

Suppose that we add one person to this population at the neutral level (of personal good) without affecting the personal good of anyone else in the population. What will the well-being of this new person be?

There are two cases to consider. The first is the case in which the level of well-being of the new individual is less than w^* . In this case, the new individual will be the only person with a complaint. Solving (11.2) for w_ν using the complaint function given in (11.3) yields:

$$w_\nu = \frac{\nu + \alpha w^*}{1 + \alpha} \quad (11.4)$$

In this first case, the neutral level of *well-being* (as opposed to personal good) depends directly on the level of well-being of the rest of the population.

The second case to consider is that in which the level of well-being of the new individual is greater than w^* . With the simplified complaint function given in (11.3), the new individual will have no complaint regarding inequality. Equation (11.2) will

¹⁹ I prefer ATBO, but here use BOP for simplicity.

reduce to:

$$w_\nu = \nu \quad (11.5)$$

That is, if we add to a population an individual whose well-being is greater than that of the members of the population – if we add an individual with *no* comparative complaint – then their well-being is equal to the neutral level.

There is a catch. For this addition to be truly neutral – for both Population A and Population A-plus-addition to have the same value – the total good of the original individuals must remain the same. This was an assumption we made above. In the non-egalitarian case presented in Section 11.2.2, this was reasonably straightforward: assuming that the new individual had no effect on the well-being of existing individuals entailed that he also had no effect on their personal good. But it is not so straightforward in the egalitarian case.

The total good of Population A, prior to the addition, was simply the sum of each individual's well-being, w^* . But this was a simplification. The total good of Population A was really the sum of each individual's personal good, which was the difference between their well-being and any comparative complaints. We were able to eliminate their comparative complaints because, in Population A, everyone was equal.

This is no longer true with the addition of the new person at a level of well-being above w^* . Now, the members of Population A are the worse-off in the new distribution. According to the complaint function assumed in (11.3), each one has a complaint against the new individual proportional to the difference between the new individual's level of well-being and w^* . If this is not somehow off-set, it will reduce the level of personal good of each individual, and thereby the total good of Population A.

The implication is that to off-set these new complaints, the well-being of the individuals in Population A must be increased. Of course, this need not be done by raising the level of well-being of *each* individual within the population. All that is required is that the total good of the members of Population A remains the same, and this entails that at least *some* members of that population experience an increase in their well-being if the addition of the new member is to be truly neutral.

Is this plausible? It should be. There is an obvious implication of the reduced equation in (11.5). To add an individual at the neutral level, but with well-being greater than that of the population to which he is added, to an even population, the existing population must have well-being that is *below* the neutral level of personal good. Put this way, the increase in the well-being of the existing population required for the addition of the new individual seems plausible.

11.4 WILL THE FLOATING NEUTRAL LEVEL MAKE A DIFFERENCE?

The result of the last section, that the neutral level of well-being is context-sensitive, is encouraging. But will it make a substantial difference?

The answer depends on the weight given to complaints, a survey of which is beyond the scope of our present discussion. There is, however, a *prima facie* reasonable condition on fixed-population value functions that may reduce the hoped-for influence of a variant neutral level of well-being. I call this condition *No Leveling Down*. It comes in two forms:

Strong No Leveling Down If in one alternative, everyone has at least the same well-being as in another, and someone has a higher well-being, then it is better.

Weak No Leveling Down If in one alternative, everyone has at least the same

well-being as in another, and someone has a higher well-being, then it is at least as good as the other alternative.

Strong No Leveling Down is the requirement that the general betterness relation be increasing in individual well-being. This places a strong limit on the weight that may be given to the disvalue of complaints, though I won't discuss the implications here.

One might doubt the Strong No Leveling Down condition, but still maintain that free increases in well-being, at the least, do not make the outcome *worse*. A population resulting from a free increase in well-being for one individual is *at least as good as* the original population. That is, leveling down the better/best-off does not make an outcome *better*. But Weak No Leveling Down does not require that leveling down make the outcome worse. Rather, it is the restriction that an increase in individual well-being cannot lead to a decrease in the betterness of an outcome.

One striking difference between the implications of the Strong and Weak versions of the No Leveling Down Condition is that the Weak No Leveling Down Condition does not entail a single neutral level for any given context. Additionally, an axiology that satisfies Weak No Leveling Down, but not Strong No Leveling Down, produces a *neutral range* for at least some contexts.

Consider 3 populations: Population A, consisting of 4 people with the same well-being 10; Population B, consisting of Population A plus one new person with well-being 10; and Population C, consisting of Population A plus one new person with well-being 15. Begin by comparing Population B and C. The Weak No Leveling Down Condition does not require that the value of a population be increasing in individual well-being, only that it not be *decreasing* in individual well-being. So Population C is at least as good as B, but it could be the case that it is also not better than B. If it is one of such cases where it is not better, then C is neither better nor worse than B. Given the existence of a value function for same number






Figure 11.2: Populations A, B, and C

cases (assumed above, Assumption 3 on pg. 238), C is equally as good as B.²⁰

Suppose that the individual in B had been added to A at a neutral level of well-being. Since p_5 is added at the neutral level of well-being, this does not affect the value of the population (from the definition of the Neutral Level of Well-being). B is neither better than nor worse than A, and, given completeness, B is equally as good as A.

By transitivity of ‘equally good as’, C is also equally as good as A. That is, the individual added to A to produce C was also added at a neutral level of well-being, albeit a different one from B.²¹

In this scenario, there exists a neutral range of well-being. The size and scope of it depend on what kind of improvements in individual well-being do not improve the distribution.²²

11.5 AN OBJECTION & A SOLUTION

The individualistic egalitarian might be poised to avoid the problems of a *constant* neutral level of well-being. But a similar objection can be raised against individualistic egalitarianism as against the neutral level’s dependence on the average: the level of well-being at which it would be better to add a person today should not depend on how well-off temporally distant people were. That is, whether it is better to add a person today should not depend on the quality of life during the Stone

²⁰ This is a case where the value of the increase in well-being for p_5 is outweighed by the disvalue of the inequality for the others in the population. I do not mean to suggest that B and C as depicted actually present such a case, but use them for illustration. Under Weak No Leveling Down, there could be two populations B and C where C is created by giving one individual in B a free increase in well-being, but where C is equally as good as B. An axiology that satisfies Weak No Leveling Down but not Strong No Leveling Down will have at least one such case.

²¹ Note that the levels of personal good of each individual in B and C would not be the same, though the total good would be.

²² How much this will help remains to be seen and depends on a fleshed out theory of complaints; there will certainly continue to be difficult trade-offs. For an interesting discussion of egalitarianism and Mere Addition, see Arrhenius (2009).

Age.²³ Unmodified individualistic egalitarianism — what I have elsewhere referred to as Temporally Neutral Egalitarianism (TNE) — has this implication. The neutral level of well-being, as just argued, depends on the well-being of the population to which an individual is added. If complaints are not restricted in scope, or heavily discounted for distance between the complaint owner and what that complaint is measured relative to, then complaints — and hence, the neutral level of well-being — could be affected by very temporally distant people.²⁴ If this was the reason that we rejected the dependence of the neutral level on the average, it seems that we cannot ignore this objection.²⁵

I do not have a good way to meet this objection on behalf of TNE. Indeed, I have

²³ This objection was first presented by Parfit in an argument against the value of an outcome depending on the Average Principle:

Whether [the birth of any child] would be bad, on the Average Principle, depends on facts about all previous lives. If the Ancient Egyptians had a very high quality of life, it is more likely to be bad to have a child now. It is more likely that this child's birth will lower the average quality of the lives that are ever lived. But research in Egyptology cannot be relevant to our decision whether to have children. (Parfit, 1984, 420)

This objection has also been discussed in Arrhenius and Bykvist (1995, 110ff.) as the need for an axiology to avoid the “Egyptology Conclusion” and in Broome (2004, 194-195) as the “objection from separability.”

A stronger version of this principle, the “Independence of the Utilities of the Dead Axiom”, has been suggested by Blackorby et al. (2005, Ch. 8). Independence of the Utilities of the Dead is stronger than avoidance of the Egyptology Conclusion in that it requires a ranking of two alternatives to be independent of the existence of *all* individuals whose lives are over, including not only the temporally distant but also the temporally close who died more recently up to and including the immediately preceding period. We can see now that, if complaints may be had against others who are temporally very close but not overlapping, egalitarianism violates this axiom when the fixed-number restriction is removed: the existence of those immediate dead will affect the neutral level of well-being, and, as result, whether it would be good to add a person to a population. As I discussed earlier (*supra* 79), complaints against the immediate dead are plausible. I also argued that, in light of this, the axiom does not provide a strong reason to reject egalitarianism. Hence, the focus in this section is on the weaker (but more troubling) Egyptology version of the objection. *cf.* Carlson (1998, 290ff.).

²⁴ A related problem is that the addition of people at certain levels of well-being could likewise affect the personal good of temporally distant people, like the cavemen. In particular, the effect that such new people would have would be to *reduce* the temporally distant people's personal good by causing them to have complaints of inequality.

²⁵ This is not the only reason to reject the dependence of the neutral level on the average. See, for example, Parfit's discussion of “How Only France Survives” and “Hell Three” which provide even stronger reasons to reject dependence on the average (Parfit, 1984, 421-422). WPE does not imply either of these conclusions.

separately argued that a similar problem arises for TNE even within fixed-number cases. As a result, I have argued in favor of developing a version of *Weighted Progressive Egalitarianism* (WPE). And I do have a good way to meet this objection on behalf of WPE — or rather, of avoiding it.

WPE, as noted²⁶, has two features: First, it restricts the scope of complaints as being against only those individuals with whom one has temporal overlap or who live at an earlier time (the Progressive Scope Restriction). Second, it discounts comparative complaints between temporally distant individuals (the Temporal Comparison Thesis).

In the case at hand, the Progressive Scope Restriction entails that the personal good of temporally distant individuals in the past is not affected by adding better off individuals at later times. As a result, the only way that the well-being of *past* individuals affects the neutral level of well-being is through a complaint on the part of the person added. That is, past individuals only affect the neutral level of well-being if they themselves were well off. Moreover, the Temporal Comparison Thesis further reduces this dependence of the neutral level of well-being on the well-being of cavemen and other temporally distant persons. As a result, the neutral level of well-being will depend directly on the well-being of the existent population, and those in the not-so-distant past, but not as strongly on the well-being of those in the very distant past.

Still, it may be objected, this does not entirely *avoid* the problem. So long as comparative complaints relative to those in the very distant past have *any* disvalue, then the neutral level still depends — however slightly — on facts about the well-being of temporally distant persons. That is, it may be objected that WPE implies the Egyptology Conclusion: “whether it would be good or bad that a person comes

²⁶*Supra* 204.

into existence, will depend on facts about the welfare and the number of the Ancient Egyptians.”²⁷

There are two ways of addressing this problem.

First, the problem might require further work in the formulation of WPE. There may be a limit to how far prior to an individual’s existence a claim can backtrack. This seems plausible: if backtracked derivative claims weaken with the degree of backtracking, it would not be surprising that there might be a degree of backtracking past which the claim ceases to exist.²⁸ If this is correct, then there will also cease to be overlapping claims between these temporally distant people, and so too for there to be any disvalue to comparisons with those temporally distant individuals who existed prior to the point where the backtracked claim ceased to exist.²⁹ Since those comparisons do not matter (would not have disvalue), those temporally distant individuals would not affect the neutral level. Whether this avoids the Egyptology Conclusion (or the Stone Age Conclusion, or the Whatever-Came-Before-the-Stone-Age Conclusion) depends on the temporal distance past which claims cannot backtrack — that is, whether those born in 2014 could have backtracked derivative claims to the resources of the Ancient Egyptians. *Prima facie*, this appears a promising line to pursue.

Alternatively, an advocate of WPE might bite the bullet and deny that the

²⁷Arrhenius and Bykvist (1995, 110). WPE also implies the *Astronomy Conclusion*: “Whether it would be good or bad that a person comes into existence, will depend on facts about the welfare and the number of the people living billions of light years from here, even though this decision will not affect these people’s welfare” (112).

I find the reasons to *avoid* the Astronomy Conclusion to be compelling, but less so than those for avoiding the Egyptology Conclusion. But this is a discussion for another day. There are independent reasons to think that if temporal distance matters to the strength of claims — and so to the disvalue of comparative complaints — then extreme spatial distance may also be morally relevant. As I noted earlier, a full exploration of this topic is warranted but beyond the scope of the present work.

²⁸Supra 7.3.1.

²⁹“Ceased” is probably not the best word in this context, but I trust my meaning is clear: for any individual, there is a time before their existence prior to which their claims cannot backtrack — that is, prior to which there are no resources to which they might have a claim.

Egyptology Conclusion *must* be avoided.³⁰ One could argue that the weakened dependence of the neutral level is enough to accommodate the intuition about the Egyptology Conclusion: it seems strange that the neutral level could depend at all on the welfare of temporally distant individuals, but the real implausibility was that it could depend just as heavily on the welfare of temporally distant individuals as temporally close ones. And this real implausibility has been successfully avoided by WPE.³¹

This result has promise. Whether it has the potential to resolve any paradoxes remains to be seen, but this promise at least provides reason to further develop WPE and consider more carefully its implications. Unfortunately, I must leave such work for another time.

11.6 CLOSING REMARKS

In Part II, I started to develop a theory of claims that I hoped would help answer the two questions which I left open at the end of Part I. Those questions were, first, whether a principle that limited the scope of complaints to comparisons with contemporaries or predecessors could be justified, and, second, whether a principle that discounted the contribution of a given comparison to an individual's complaint for the temporal distance between the individuals so compared could be justified.

³⁰ *cf.* Carlson (1998, 290-294). Carlson argues against the need for avoiding the Egyptology Conclusion on two grounds, one of which is relevant here: he argues that, as an axiological matter, if one accepts independence of common temporal parts of a history, one should also accept independence of common spatial parts, with the result that to accept the Independence of the Utilities of the Dead without “drawing arbitrary temporal or spatial boundaries” requires acceptance of the *Invariance of Contributive Value* (ICV). He defines ICV as the principle that “The contributive value of any life, at a given level of welfare, is the same for all populations which include this life” (291). He then argues against ICV, with the result that he rejects the Independence of the Utilities of the Dead.

Egalitarianism is also incompatible with ICV. This second method of confronting the Egyptology problem under WPE also appears to avoid problems of arbitrary line-drawing.

³¹ Unlike under WPE, the neutral level under the Average Principle and TNE both depend as heavily on the welfare of temporally distant individuals as temporally close ones.

I also said that the fate of telic egalitarianism (of some stripe) depended on the availability of such justification.

My limited conclusion from Part III is that such an approach shows promise. As may by now be clear, the theory requires development: to answer questions that are hanging around the edges, and to shore up some of the arguments by finding better reasons for accepting their assumptions. But one thing remains clear to me, as it had in the beginning: considering how time — and, more accurately, temporal distance — makes a morally relevant difference has the potential to shed new light on old problems.

Bibliography

Anscombe, G. E. M. (1958). Modern Moral Philosophy. *Philosophy* 33(124), 1–19.

Arrhenius, G. (2000). *Future Generations: A Challenge for Moral Theory*. FD-Diss., Uppsala University, Uppsala.

Arrhenius, G. (2003). Feldman’s Desert-Adjusted Utilitarianism and Population Ethics. *Utilitas* 15(2), 225–36.

Arrhenius, G. (2006). Desert as Fit: An Axiomatic Analysis. In K. McDaniel, R. Feldman, J. R. Raibley, and M. J. Zimmerman (Eds.), *The Good, the Right, Life, and Death: Essays in Honor of Fred Feldman*, 3–18. Aldershot: Ashgate.

Arrhenius, G. (2007). Meritarian Axiologies and Distributive Justice. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, and D. Egonsson (Eds.), *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. www.fil.lu.se/hommageawlodek.

Arrhenius, G. (2009). Egalitarianism and Population Change. In A. Gosseries and L. H. Meyer (Eds.), *Intergenerational Justice*, 323–46. Oxford: Oxford University Press.

Arrhenius, G. and K. Bykvist (1995). *Future Generations and Interpersonal Compensations: Moral Aspects of Energy Use*. Uppsala Prints and Preprints in Philosophy. No. 21. Uppsala: Department of Philosophy, University of Uppsala.

- Barro, R. J., X. Sala-i Martin, and N. Inc. (2004). *Economic Growth* (2nd ed.). Cambridge, Mass.: MIT Press.
- Barry, B. and R. I. Sikora (1978). *Obligations to Future Generations*. Philadelphia: Temple University Press.
- Blackburn, S. (1996). *The Oxford Dictionary of Philosophy*. Oxford: Oxford University Press.
- Blackorby, C., W. Bossert, and D. Donaldson (1995). Intertemporal Population Ethics: Critical-Level Utilitarian Principles. *Econometrica* 63(6), 1303–20.
- Blackorby, C., W. Bossert, and D. Donaldson (2005). *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press.
- Broome, J. (1984a). Selecting People Randomly. *Ethics* 95(1), 38–55.
- Broome, J. (1984b). Uncertainty and Fairness. *The Economic Journal* 94(375), 624–32.
- Broome, J. (1989). What's the Good of Equality. In J. Hey (Ed.), *Current Issues in Microeconomics*, 236–62. Macmillan.
- Broome, J. (1991). *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Blackwell.
- Broome, J. (1992). *Counting the Cost of Global Warming*. Cambridge: White Horse Press. A Report to the Economic and Social Research Council on Research by John Broome and David Ulph.
- Broome, J. (1999). *Ethics Out of Economics*. Cambridge: Cambridge University Press.

- Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press.
- Carlson, E. (1998). Mere Addition and Two Trilemmas of Population Ethics. *Economics and Philosophy* 14, 283–306.
- Clayton, M. and A. Williams (2000). *The Ideal of Equality*. Basingstoke: Palgrave Macmillan. 2002.
- Cowen, T. (1992). Consequentialism Implies a Zero Rate of Intergenerational Discount. In P. Laslett and J. S. Fishkin (Eds.), *Justice Between Age Groups and Generations*, Philosophy, Politics, and Society, Sixth Series, 162–68. New Haven: Yale University Press.
- Cowen, T. and D. Parfit (1992). Against the Social Discount Rate. In P. Laslett and J. S. Fishkin (Eds.), *Justice Between Age Groups and Generations*, Philosophy, Politics, and Society, Sixth Series, 144–61. New Haven: Yale University Press.
- Dalton, H. (1920). The Measurement of the Inequality of Incomes. *The Economic Journal* 30(119), 348–61.
- Dummett, M. (1964). Bringing About the Past. *The Philosophical Review* 73(3), 338–59.
- Feldman, F. (1995a). Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice. *Philosophy and Phenomenological Research* 55(3), 567–85. Reprinted in Feldman (1997).
- Feldman, F. (1995b). Desert: Reconsideration of Some Received Wisdom. *Mind* 104(413), 63–77. Reprinted in Feldman (1997).
- Feldman, F. (1995c). Justice, Desert, and the Repugnant Conclusion. *Utilitas* 7(2), 189–206. Reprinted in Feldman (1997).

- Feldman, F. (1997). *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy*. Cambridge Studies in Philosophy. Cambridge: Cambridge University Press.
- Griffin, J. (1986). *Well-being: Its Meaning, Measurement and Moral Importance*. Oxford: Oxford University Press.
- Griffin, J. (2008). *On Human Rights*. Oxford: Oxford University Press.
- Hawthorne, J. (2006). Three-Dimensionalism. In J. Hawthorne (Ed.), *Metaphysical Essays*, 85–110. Oxford: Oxford University Press.
- Hume, D. (1978). *A Treatise of Human Nature* (2nd ed.). Oxford: Oxford University Press.
- Hurka, T. (2001). The Common Structure of Virtue and Desert. *Ethics* 112(1), 6–31.
- Hurka, T. (2003). Desert: Individualistic and Holistic. In S. Olsaretti (Ed.), *Desert and Justice*, 45–68. Oxford: Oxford University Press.
- Kagan, S. (1998). *Normative Ethics*. Dimensions of Philosophy Series. Boulder, Colo.: Westview Press.
- Kagan, S. (1999). Equality and Desert. In L. P. Pojman and O. McLeod (Eds.), *What Do We Deserve?: A Reader on Justice and Desert*, 298–314. Oxford: Oxford University Press.
- Kagan, S. (2003). Comparative Desert. In S. Olsaretti (Ed.), *Desert and Justice*, 93–122. Oxford: Oxford University Press.
- Kagan, S. (Forthcoming). *The Geometry of Desert*. Oxford: Oxford University Press.

- Kahneman, D., B. L. Fredrickson, C. A. Schreiber, and D. A. Redelmeier (1993). When More Pain Is Preferred to Less: Adding a Better End. *Psychological Science* 4(6), 401–405.
- Kamm, F. M. (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford Ethics Series. Oxford: Oxford University Press.
- Kleinig, J. (1971). The Concept of Desert. *American Philosophical Quarterly* 71(8), 71–78.
- Laslett, P. and J. S. Fishkin (1992). *Justice Between Age Groups and Generations*. Philosophy, Politics, and Society, Sixth Series. New Haven: Yale University Press.
- Lewis, D. (1986a). Causation (with Postscripts). In D. Lewis (Ed.), *Philosophical Papers*, Volume II, 159–213. Oxford: Oxford University Press.
- Lewis, D. (1986b). Counterfactual Dependence and Time’s Arrow (with Postscripts). In D. Lewis (Ed.), *Philosophical Papers*, Volume II, 32–66. Oxford: Oxford University Press.
- Lewis, D. (1986c). The Paradoxes of Time Travel. In D. Lewis (Ed.), *Philosophical Papers*, Volume II, 67–80. Oxford: Oxford University Press.
- McKerlie, D. (1989). Equality and Time. *Ethics* 99(3), 475–91.
- Nagel, T. (1970). *The Possibility of Altruism*. Princeton: Princeton University Press.
- Nagel, T. (1979). *Mortal Questions*. Cambridge: Cambridge University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press. Reprinted in 1987 (with corrections).

- Parfit, D. (1995). *Equality or Priority?* Lindley Lecture. Lawrence, Kansas: Dept. of Philosophy, University of Kansas. Reprinted in Clayton, M. and A. Williams (2000).
- Pigou, A. C. (1912). *Wealth and Welfare*. London: Macmillan.
- Rawls, J. (1999). *A Theory of Justice* (Revised ed.). Cambridge, Mass.: Harvard University Press.
- Robbins, T. (2010). Arizona Budget Cuts Put Organ Transplants At Risk. *NPR News (online)* (17 November 2010). <http://www.npr.org/2010/11/11/131215308/arizona-budget-cuts-put-organ-transplants-at-risk>.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge, Mass.: Harvard University Press.
- Sen, A. (1973). *On Economic Inequality*. Radcliffe Lectures. Oxford: Oxford University Press.
- Sider, T. (2001). *Four-Dimensionalism : An Ontology of Persistence and Time*. Oxford: Oxford University Press.
- Strawson, P. (2003). Freedom and Resentment. In G. Watson (Ed.), *Free Will* (2nd ed.), 72–93. Oxford: Oxford University Press.
- Temkin, L. S. (1993). *Inequality*. Oxford Ethics Series. Oxford: Oxford University Press.
- Watson, G. (2003). *Free Will* (2nd ed.). Oxford Readings in Philosophy. Oxford: Oxford University Press.
- Wiggins, D. (1998). Claims of Need. In D. Wiggins (Ed.), *Needs, Values, Truth: Essays in the Philosophy of Value* (3rd ed.), 1–58. Oxford: Oxford University Press.

-
- Zimmerman, M. J. (2002). Taking Luck Seriously. *Journal of Philosophy* 99(11), 553–76.

Part IV

Appendices

MUSHROOM GROWTH MODEL

The following is a discussion of the growth model presented in Chapter 6.

A.1 SIMPLE MUSHROOM GROWTH MODEL (CONTINUOUS TIME)

A.1.1 SIMPLE MUSHROOM GROWTH MODEL: WITHOUT INTERFERENCE

Suppose that there is a world in which there is only one resource, a magic mushroom. The mushroom grows continuously. The mushroom's rate of growth, r , is given as:

$$r = \text{growth rate} = \frac{\dot{M}(t)}{M(t)} \quad (\text{A.1})$$

where $M(t)$ is the size of the mushroom at time t , and $\dot{M}(t)$ is the instantaneous rate of change of the mushroom size.¹ The growth formula (without consumption)

¹ I follow the convention of using a single dot to denote the first derivative ($\dot{M}(t) \equiv \frac{dM}{dt}$).

is thus given as:

$$\dot{M}(t) = rM(t) \quad (\text{A.2})$$

The solution to this differential equation is:

$$M(t) = M_0 e^{rt} \quad (\text{A.3})$$

To confirm this, note that the first derivative of $M(t)$ is:

$$\dot{M}(t) = r \cdot M_0 e^{rt} \quad (\text{A.4})$$

which is equivalent to equation (A.2).

A.1.2 SIMPLE MUSHROOM GROWTH MODEL: WITH CONSUMPTION

But Equation (A.3) will not provide the correct result of mushroom resource at a given time in our example since the amount of mushroom at a given time also depends on how much of the mushroom stock is consumed. The growth formula with consumption is the growth formula without consumption less the amount consumed at a time:

$$\dot{M}(t) = rM(t) - C(t) \quad (\text{A.5})$$

where $C(t)$ is the amount consumed at t , r is the growth rate of the mushroom (without consumption), $M(t)$ is the size of the mushroom at time t , and $\dot{M}(t)$ is the instantaneous rate of change of the mushroom size. The solution to Equation (A.5) will be the amount of mushroom that would have grown without consumption ($M_0 e^{rt}$) less the amount that was consumed and what would have grown from that

consumption $(\int_0^t c(\tau)e^{r(t-\tau)} d\tau)$:

$$M(t) = M_0 e^{rt} - \int_0^t C(\tau) e^{r(t-\tau)} d\tau \quad (\text{A.6})$$

To confirm this, note that the first derivative of $M(t)$ is:

$$\dot{M}(t) = r \left(M_0 e^{rt} - \int_0^t C(\tau) e^{r(t-\tau)} d\tau \right) - C(t) \quad (\text{A.7})$$

which is equivalent to Equation (A.5).²

A.1.3 INDIVIDUAL LEVELS OF CONSUMPTION

The chapter to which this model relates, Chapter 6, is concerned with individual claims to the mushroom, and, more particularly, to consuming the mushroom. How can our model accommodate this?

The consumption at a time, $C(t)$, is the sum of individual consumption at a time, and is constrained by the amount of mushroom available at that time:

$$C(t) = \sum_{i=1}^k \sum_{j=1}^n C_{i,j}(t) \quad \text{where } 0 \leq C(t) \leq M(t) \quad (\text{A.8})$$

Here, $C_{i,j}(t)$ denotes the consumption at time t of person $p_{i,j}$.

In the model developed in Chapter 6, not everyone lives at each time. But the equation (??) can accommodate this: $C_{i,j}(t)$ will only have a positive value for individuals who consume at time t . Those persons, $p_{i,j}$, who do not exist at time t , will not be able to consume the mushroom at time t . Hence, $C_{i,j}(t) = 0$ when $p_{i,j}$ does not exist.

There is then a further simplifying feature that arises from the non-overlapping of generations assumed in most of our examples. We may take the total consumption

² Many thanks to John Broome for his help with these.

at a given time to be the sum of the consumption of those individuals who are members of the generation then existing—that is, to be the consumption of the generation then existing (and only that generation) at that time:

$$C(t) = C_i(t) \tag{A.9}$$

$$= \sum_{j=1}^n C_{i,j}(t) \tag{A.10}$$

where i is the generation alive at t , and $C(t)$ is subject to the same constraint as before ($0 \leq C(t) \leq M(t)$).

A.1.4 CLAIMS TO THE AGGREGATE

In 6.3³, I supposed that each person had a claim to an *equal share* of the historic mushroom aggregate. But how should we “count” this aggregate?

For simplicity, I assumed that pieces of mushroom are historically extended. Since we’re interested in *consumable* mushroom, the mushroom should be counted at point of consumption (whether or not it is actually “consumed”, or just destroyed). As a result, the historic mushroom aggregate may be give as the total consumption over time:

$$\int_0^T C(t) dt \tag{A.11}$$

where T is the last time of consumption (i.e. the time after which there is no longer any mushroom), and $C(t)$ is the consumption at a time.

That is, the size of mushroom to which each individual has a claim, denoted $m_{i,j}$, is given by the historic mushroom aggregate divided by the historic population:

$$m_{i,j} = \frac{\int_0^T C(t) dt}{k \cdot n} \tag{A.12}$$

³p. 128

$C(t)$ is itself constrained by $M(t)$. This gives us boundary conditions:

$$0 \leq C(t) \leq M(t) \quad (\text{A.13})$$

Since the amount which *can* be consumed is constrained by the amount of mushroom, which itself is dependent on the growth rate and levels of consumption ((A.6)), it is clear that the size of an individual's claim to the mushroom will depend on how much is consumed and when.

If all claims are satisfied on this view, then everyone will consume the same amount:

$$C_{i,j} = m \quad \text{for all } i, j \quad (\text{A.14})$$

Since it has been assumed that everyone lives for the same amount of time, we may also assume, for simplicity, that everyone has the same pattern of consumption, and consumes the same amount at each time they are alive. That is, assume:

$$C_{i,j}(t) = m \quad \text{for } t_{g_i \text{ born}} \leq t \leq t_{g_i \text{ dies}} \quad (\text{A.15})$$

$$= 0 \quad \text{for all other times} \quad (\text{A.16})$$

where m is a constant. As a result:

$$C_{i,j} = \int_{t_{g_i \text{ born}}}^{t_{g_i \text{ dies}}} m \, dt \quad (\text{A.17})$$

Since only one generation exists at any given time, each with the same generational population of n , and since there are no gaps between generations, we can rewrite $C(t)$ as:

$$C(t) = n \cdot m \quad (\text{A.18})$$

Combining (A.12), (A.14), and (A.18), we find that the amount claimed, and con-

sumed, by an individual in this history is:

$$m_{i,j} = \frac{\int_0^T n \cdot m \, dt}{k \cdot n} \quad (\text{A.19})$$

This is equivalent to:

$$m_{i,j} = \frac{(n \cdot m) \int_0^T dt}{k \cdot n} \quad (\text{A.20})$$

$$m_{i,j} = \frac{(n \cdot m)(T - 0)}{k \cdot n} \quad (\text{A.21})$$

$$m_{i,j} = \frac{mT}{k} \quad (\text{A.22})$$

What is the value of m ? Assuming that, in this history, the entire mushroom is consumed, then we know that:

$$C(T) = M(T) \quad (\text{A.23})$$

That is, all of the remaining mushroom is consumed at the end at time T . This may be rewritten in terms of m :

$$n \cdot m = M(T) \quad (\text{A.24})$$

$$m = \frac{M(T)}{n} \quad (\text{A.25})$$

$$m = \frac{M_0 e^{rT} - \int_0^T C(\tau) e^{r(T-\tau)} \, d\tau}{n} \quad (\text{A.26})$$

$$m = \frac{M_0 e^{rT} - \int_0^T n m e^{r(T-\tau)} \, d\tau}{n} \quad (\text{A.27})$$

$$m = \frac{M_0 e^{rT} - n m \left(\frac{1 - e^{rT}}{-r} \right)}{n} \quad (\text{A.28})$$

$$nm \left(1 - \frac{1 - e^{rT}}{r} \right) = M_0 e^{rT} \quad (\text{A.29})$$

$$m = \frac{r M_0 e^{rT}}{n(r + e^{rT} - 1)} \quad (\text{A.30})$$

MUSHROOM BLIGHT

In 6.3, there is an unexpected mushroom blight at t_{blight} . As an example, suppose that the blight destroys a part of the mushroom of size m_{blight} . Prior to the mushroom blight, the amount of mushroom will be given as above ---:

$$M(t) = M_0 e^{rt} - \int_0^t C(\tau) e^{r(t-\tau)} d\tau \quad \text{for } 0 \leq t \leq t_{\text{blight}} \quad (\text{A.31})$$

After the blight, however, the amount of mushroom at time t has been reduced by the amount destroyed in the blight, as well as whatever would have grown from that portion:

$$M(t) = M_0 e^{rt} - \int_0^t C(\tau) e^{r(t-\tau)} d\tau - m_{\text{blight}} e^{r(t-t_{\text{blight}})} \quad \text{for } t_{\text{blight}} \leq t \leq T \quad (\text{A.32})$$

Note that, given the blight, it is no longer the case that each individual may consume m at each time during their life. To see this, suspend belief and suppose that they did — that negative consumption is possible. Then the amount of mushroom remaining at T can be determined as follows:

$$M(T) = M_0 e^{rT} - \int_0^T n m e^{r(T-\tau)} d\tau - m_{\text{blight}} e^{r(T-t_{\text{blight}})} \quad (\text{A.33})$$

$$= M_0 e^{rT} - n m \left(\frac{1 - e^{rT}}{-r} \right) - m_{\text{blight}} e^{r(T-t_{\text{blight}})} \quad (\text{A.34})$$

$$= M_0 e^{rT} - n \left(\frac{r M_0 e^{rT}}{n(r + e^{rT} - 1)} \right) \left(\frac{1 - e^{rT}}{-r} \right) - m_{\text{blight}} e^{r(T-t_{\text{blight}})} \quad (\text{A.35})$$

$$= \frac{r M_0 e^{rT} - m_{\text{blight}} e^{r(T-t_{\text{blight}})} (r + e^{rT} - 1)}{(r + e^{rT} - 1)} \quad (\text{A.36})$$

$$= \left(\frac{r M_0 e^{rT}}{(r + e^{rT} - 1)} \right) - m_{\text{blight}} e^{r(T-t_{\text{blight}})} \quad (\text{A.37})$$

$$= (n \cdot m) - m_{\text{blight}} e^{r(T-t_{\text{blight}})} \quad (\text{A.38})$$

As we can see, there is less than enough remaining at T for each of the n people to consume m at that time.