

The Baseline Problem in Tax Analysis

by

Ryan Tonkin

B.A., University of Victoria, 2010

M.A., University of Victoria, 2011

J.D., Harvard Law School, 2014

A Dissertation Submitted in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Philosophy

©Ryan Tonkin, 2025  
University of Victoria

All rights reserved. This dissertation may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

*We acknowledge and respect the Ləkʷəŋən (Songhees and Xʷsepsəm/Esquimalt) Peoples on whose territory the university stands, and the Ləkʷəŋən and WSÁNEĆ Peoples whose historical relationships with the land continue to this day.*

The Baseline Problem in Tax Analysis

by

Ryan Tonkin

B.A., University of Victoria, 2010

M.A., University of Victoria, 2011

J.D., Harvard Law School, 2014

Supervisory Committee

Dr. Colin Macleod, Supervisor  
Department of Philosophy

Dr. Peter Dietsch, Departmental Member  
Department of Philosophy

Dr. Bradley Bryan, Outside Member  
Department of Law

## **Abstract**

The purpose of this project is to articulate and defend the view that both fairness and efficiency as they are understood in a great deal of tax analysis are plagued by a similar kind of comparative conceptual error, which I characterize as a baseline problem. I begin with consideration of the notions of property and ownership that best capture our existing legal and institutional practice. This practice reveals how ostensibly pretax descriptions of resources are institutionally constructed and highly contingent. I then turn to tax burden analysis as it is understood in tax policy and the orthodox approach to public finance. I join a handful of scholars who argue that this orthodoxy commits a serious conceptual error. As an alternative, I recommend conceiving of taxation as a technology embedded in a comprehensive social architecture. Resistance to this kind of approach is often directed at its alleged implausible commitment to a thoroughly conventional conception of property rights. However, I argue it has no such commitment and is consistent with a strong commitment to economic liberties. I then argue that a parallel baseline problem can be identified in two economic concepts routinely applied to assess efficiency in tax policy: Pareto efficiency and deadweight loss. I argue that both depend on comparative baselines that are often implausible or unjustified, and at the same time ignored or obscured. Finally, I consider the practice of reporting tax expenditures and the public discourse about “fair shares” as examples of baseline problems in practice.

## Table of Contents

Supervisory Committee .....	ii
Abstract .....	iii
Table of Contents .....	iv
List of Tables .....	vii
List of Figures .....	viii
Acknowledgments .....	ix
Dedication .....	x
Project Background and Motivation .....	1
Project Overview .....	9
Chapter 1: Mine and Thine .....	14
1.1 Chapter 1 Introduction .....	14
1.2 The Property Strategy .....	17
1.3 Rights or Conventions .....	21
1.4 Conceptions of Property and Ownership .....	26
1.4.1 Foundations .....	26
1.4.2 Relations with Things .....	28
1.4.3 Relations with Persons .....	30
1.4.4 The Gaian Inconsistency .....	32
1.4.5 Revisiting the Triadic Formulation .....	36
1.4.6 The Disintegration of Dominion .....	39
1.4.7 Disintegration vs. Unified Core .....	44
1.4.8 Limitations of the Legal Paradigm .....	50
1.4.9 The Objects of Ownership: from Self to Wealth .....	55
1.5 Chapter 1 Conclusion .....	62
Chapter 2: Normative Taxation .....	66
2.1 Chapter 2 Introduction .....	66
2.2 Taxation as a Normative Inquiry .....	67
2.3 Tax Burden Analysis .....	71
2.3.1 Horizontal and Vertical Equity .....	72
2.3.2 Ability to Pay .....	74
2.3.2.1 Endowment .....	74
2.3.2.2 Equal Sacrifice .....	76
2.3.3 The Benefit Principle .....	78
2.3.4 Tax Constitutionalism .....	81
2.4 Tax as Social Architecture .....	83
2.4.1 Constructing Entitlements .....	84
2.4.2 Myopia and Holism .....	91
2.4.3 Optimal Tax .....	93
2.4.4 Choice Architecture .....	95
2.4.4.1 Public Goods .....	99
2.5 Emerging Issues .....	102
2.5.1 Monetary Theory .....	102
2.5.2 Tax Among Nations .....	103
2.6 Chapter 2 Conclusion .....	106

Chapter 3: The Baseline Problem in Tax Fairness.....	108
3.1 Chapter 3 Introduction .....	108
3.2 Taxation and the Property Constraint .....	109
3.2.1 The Conceptual Argument.....	112
3.2.2 The Incoherence Thesis .....	114
3.2.3 Property Rights and Legal Conventionalism.....	116
3.2.4 Conventions .....	119
3.2.5 The Property-Law Unity .....	125
3.2.6 Economic Exceptionalism .....	129
3.2.7 Implications of Exceptionalism .....	136
3.2.8 Moral Property Constraints on Taxation.....	141
3.3 Chapter 3 Conclusion.....	146
Chapter 4: The Baseline Problem in Tax Efficiency .....	148
4.1 Chapter 4 Introduction .....	148
4.2 The Pareto Principle.....	150
4.2.1 Introduction.....	150
4.2.2 Pareto Concepts .....	151
4.2.3 Neutrality in Public Policy.....	153
4.2.4 An Example Choice .....	155
4.2.5 Comparative Baselines.....	156
4.2.6 Uncontroversial Conflation.....	159
4.2.7 Pareto Principle Conclusion.....	163
4.3 Deadweight Loss.....	163
4.3.1 Introduction.....	163
4.3.2 Taxing the Oil Market.....	166
4.3.3 Surplus Value.....	168
4.3.4 The Tax Wedge.....	170
4.3.5 The “Market” Baseline .....	172
4.3.6 Deadweight Loss Conclusion .....	175
4.4 Chapter 4 Conclusion.....	176
Chapter 5: The Baseline Problem in Practice .....	177
5.1 Chapter 5 Introduction .....	177
5.2 Tax Expenditures .....	181
5.2.1 Introduction.....	181
5.2.2 Spending by Not Taxing .....	182
5.2.3 Spending in the Shadows.....	184
5.2.4 The Tax Expenditure Enigma .....	187
5.2.5 The ‘Last Penny’ .....	190
5.2.6 The Normal Tax.....	192
5.2.7 The Canadian Framework.....	195
5.2.7.1 Canada’s Flawed Baseline .....	197
5.2.7.2 Counterfactuals and Nominal Revenue .....	204
5.2.8 The JCT Reform Proposal .....	208
5.2.8.1 Perceived Neutrality.....	209
5.2.8.2 Cryptonormativity in ‘Present Law’ .....	211
5.2.8.3 Cryptonormativity in ‘Efficiency’ .....	214

5.2.9 Candid Normativity .....	223
5.3 Taxation and ‘Fair Shares’ .....	227
5.4 Conclusion .....	238
5.5 Chapter 5 Conclusion.....	238
Project Conclusion.....	241
References.....	245

## List of Tables

Table 1: The Gaian Inconsistency.....	34
Table 2: A Choice Among Three Tax Schedules .....	156
Table 3: A Choice Among Three Tax Schedules (repeated) .....	157
Table 4: Pareto Comparisons of Three Tax Schedules .....	158
Table 5: A Choice Among Four Tax Schedules .....	160
Table 6: Pareto Comparisons of Four Tax Schedules.....	161

## List of Figures

Figure 1: Oil Market with Tax on Producers .....	167
Figure 2: Oil Market with Tax on Consumers .....	168
Figure 3: Market Surplus Without Tax.....	169
Figure 4: Modeling Tax as Deadweight Loss .....	171

## Acknowledgments

My sincere thanks first to the members of my committee. This project would not have been possible without your shared expertise in philosophy, law, and economics. But more than that, each of you supported both me and this project in ways that go well beyond the obligations of a committee member. Thanks to Colin for your patience, wisdom, and commitment to what matters. It has been almost two decades, and no one has taught me more about what it means to be a philosopher. Thanks to Peter for interrogating every detail while never losing sight of the bigger picture: practice is ever the final test of theory. I have always admired your ability to be at once deeply incisive and relentlessly constructive. Thanks to Brad for bringing me aboard to study Canadian tax law and, in particular, its interface with Indigenous people. Our many conversations lie behind several key arguments in this project. I genuinely cannot imagine a better team.

I also want to acknowledge the time and contributions of the many other faculty members who supported this project. Thanks in particular to Audrey Yap for expertly separating and then weaving the key threads, and to Carrie Klatt for making my University of Victoria journey possible.

In addition to faculty members, many of my peers in the undergraduate and graduate programs took the time to discuss my unusual topic. You know who you are; thanks to you all.

I also want to acknowledge the University of Victoria Philosophy Department as a whole for its flexibility and understanding as I progressed (and at times regressed) through this program. A PhD is long, and my path saw life, death, sickness, health, and more in between. I was supported throughout.

I am grateful to the many sources from which I received financial support during my program. In particular, I am grateful to the Pierre Elliot Trudeau Foundation, Vanier Canada Graduate Scholarships, BC's Office of the Human Rights Commissioner, and the University of Victoria.

Finally, thanks to everyone in my personal life who supported me all these years. Thanks in particular to Nathan for dedicating his time and mind on numerous occasions, and to Carolyn for helping to keep the project on track.

## Dedication

For Amanda and Annabelle, who are effortlessly the best people I know.

*“I spoke about wings, you just flew.”*



*[Tax Justice is] not a matter of applying  
some equitable-seeming function  
to a morally arbitrary initial distribution.*

Liam Murphy & Thomas Nagel

## Project Background and Motivation

In 1879 John Stuart Mill lamented that the great majority were no longer enslaved by force of law and yet remained bound by chains of poverty.

No longer enslaved or made dependent by force of law, the great majority are so by force of poverty; they are chained to a place, to an occupation, and to conformity with the will of an employer, and debarred by the accident of birth both from the enjoyments, and from the mental and moral advantages, which others inherit without exertion and independently of desert. That this is an evil equal to almost any of those against which mankind have hitherto struggled, the poor are not wrong in believing. Is it a necessary evil? They are told so by, those who do not feel it---by those who have gained the prizes in the lottery of life. But it was also said that slavery, that despotism, that all the privileges of oligarchy were necessary (Mill, 1879, 10-21).

Nearly a century and a half later, even the most prosperous nations remain riven by poverty, substantially reducing the life prospects of the great majority. As Mill implies, poverty is not a natural disaster, but rather an unnecessary social evil. That it exists against a background of abundance is the worst manifestation of rampant inequality. Yet now more than ever there is reason to doubt Mill's separation of the forces of law and poverty. As I will argue throughout this project, law is a powerful force in constructing economic outcomes.<sup>1</sup>

Because inequality and its consequences are an expression of human agency, there is need for justification. Indeed, Thomas Piketty begins his *Capital and Ideology* with the observation that “[e]very human society must justify its inequalities: unless reasons for them are found, the whole political and social edifice stands in danger of collapse” (2020, p. 1). In his study he finds justificatory themes of social stability, individual desert, overall prosperity, divine command, and so on. Each society “develops a range of contradictory discourses and ideologies

---

<sup>1</sup> Rousseau wrote that “Under bad governments, [moral] equality is only apparent and illusory: all it does is to keep the pauper in his poverty and the rich man in the position he has usurped. Laws in fact are always useful to those who have possessions and harmful to those who don't; from which it follows that the social state is advantageous to men only when everyone has something and no-one has too much” (1.9, fn. 3). A similar sentiment is expressed in Anatole France's famous 1894 observation that “In its majestic equality, the law forbids rich and poor alike to sleep under bridges, beg in the streets and steal loaves of bread.”

for the purpose of legitimizing the inequality that already exists or that people believe should exist” (Piketty, 2020, p. 1). The social, economic, and political clash of these discourses gives rise to a dominant justificatory narrative. The dominant narrative always includes a theory of borders and a theory of property. A theory of borders is an account of who belongs to the political community. It assigns (and denies) what Hannah Arendt called “the right to have rights” (1951, p. 296). It is also an account of what territory the society governs, its institutions, and its relations to other political communities. A theory of property describes the forms of ownership, as well as the legal and other procedures that regulate property relations. At the core of every justificatory narrative, therefore, we will find a story about who can own what, and how.

Central to Piketty’s analysis is his resolution that “many paths are possible” (2020, p. 7). He rejects the view that individuals are midwives to an “old society pregnant with a new one” (Marx, 1887, p. 534). History reveals that the transition from feudalism to capitalism, for example, did not occur as a mechanical and definite response to the industrial revolution. Instead, we find myriad and highly varied social, political, and economic systems. Nothing is determined, and ideas matter in shaping the course of history. Indeed, they may matter most when, like water to fish, they are so entrenched that they barely register as ideology at all.

“...often it is those who claim to be purely pragmatic who are in fact most ‘ideological’ (in the pejorative sense): their claim to be post-ideological barely conceals their disdain for evidence, historical ignorance, distorting biases, and class interests” (2020, p. 9).

This echoes Keynes’ earlier observation about the power of ideas, especially when the status quo masquerades as free of intellectual influences.

“...the ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than is commonly understood. Indeed the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influences, are usually the slaves of some defunct economist” (Keynes, 1936, p. 234).

Egalitarian successes in the mid-twentieth century are a case in point, as they arose from a seizure of power that “was ideological and intellectual before it was political” and thus political transformation was “won above all on the battleground of ideas” (Piketty, 2020, p. 36).

Inequality is neither economic nor technological; it is ideological and political . . . In other words, the market and competition, profits and wages, capital and debt, skilled and unskilled workers, natives and aliens, tax havens and competitiveness—none of these things exist as such. All are social and historical constructs, which depend entirely on the legal, fiscal, education, and political systems that people choose to adopt and the conceptual definitions they choose to work with (Piketty, 2020, p. 7).

Given the power of ideas and a world of “many paths,” it is not enough to proceed destructively without a clear sense of alternatives.<sup>2</sup> As Marx warned, “the philosophers have only interpreted the world, in various ways; the point is to change it” (1888/1969, p. 15). The consequences of directionless revolution may be “just as dangerous as the quasi-sacred property relations that the revolution sought to overthrow” (Piketty, 2020, p. 8). This was especially clear when extreme inequality in France returned within only decades after the fall of the Ancien Régime (Morrison & Snyder, 2000).

A clear statement of alternatives is a bulwark against those who try “to naturalize inequality or deny the existence of alternative forms of social organization” (Piketty, 2020, p. 12). As Mill earlier warned, human beings “tend to consecrate existing facts, and to declare what does not yet exist, either pernicious or impracticable” (1880, 21). Of course, that the ideological foundations of inequality are constructed does not imply that their transformation is possible “by a wave of some magic wand” (2020, p. 12). Instead, Piketty maintains that persuasive alternatives require clear conceptual underpinnings. We must “choose the right terminology and

---

<sup>2</sup> Jonathan Wolff (2002, p.3) humorously relates that “one demonstrator at a recent anti-capitalist demonstration held up a banner reading ‘Replace Capitalism with Something Nice’.”

conceptual framework” (Piketty, 2020, p. 24). In particular, no where in the world “has enough attention been paid to transcending private property in its present form” (Piketty, 2020, p. 34).

I started this project to better understand the role of tax analysis in our own “inequality regime” (Piketty, 2020) and so in realizing (or frustrating) distributive justice. I discovered that many people, mostly economists and lawyers, have made serious efforts to develop principles of fairness and efficiency that are widely thought of as the normative core of justice in taxation. These principles generally respond to the “orthodox” public finance approach to tax reform and design, which goes something like this:

We begin with the requirement that we must raise a given, and fixed, amount of revenue required for public activity. This revenue can be simply thought of as a pile of treasure to be obtained. Then the question is: how should the cost of accumulating this pile be spread across the taxpaying population so that, first, the distribution of the cost across the relevant individuals is *fair*; and so that second, the aggregate cost to the taxpaying population is minimised (that is, the tax is as *efficient* as possible) (Brennan, 1984, p. 83).

Responses vary. For example, some people argue that the cost should be borne in proportion to benefit, since not everyone benefits equally from the taxing and spending of the state. Others argue that people with roughly the same amount of treasure should contribute roughly the same amount to the public revenue, while those with more should contribute more. As noted, many of the underlying arguments were developed and refined by economists and lawyers rather than philosophers. Given the centrality of normative principles to resolving this sort of problem, it is somewhat surprising that philosophers today offer little response. However, what philosophers have said is, I think, the most interesting. Instead of responding to this problem directly, the most influential work of modern philosophy on taxation—Liam Murphy and Thomas Nagel’s 2002 book, *The Myth of Ownership*—rejects the orthodox paradigm entirely. Instead, their central claims are that (1) property is conventional, (2) entitlement to a pretax distribution of resources is

incoherent, and (3) it is myopic to evaluate tax justice as distinct from the overall effect of social institutions on people's lives. The heart of their response is that taxpayer entitlements to "treasure" cannot exist prior to justice in taxation, and so it makes no sense to fairly and efficiently distribute the cost across such nonexistent entitlements.

To say that *The Myth of Ownership* has been highly influential is not to say that it has been widely persuasive. While one cannot write about tax justice without at least mentioning the text, the literature reveals few who accept its central claims. Perhaps this is because tax policy analysis has largely moved on from tax burden analysis (which assumes entitlement to a pretax distribution of resources in just the way Murphy and Nagel want to deny), at least in economics.<sup>3</sup> In its stead, the "incentive compatibility approach" (optimal tax) is apparently no longer concerned with archaic notions like horizontal and vertical equity meant to guide fairness in distributing the costs of taxation. However, I do not think this explanation is fully satisfying. The dominance of optimal tax modeling is not absolute, and nor is it universally adopted outside of academic economics. Indeed, Canada's first Parliamentary Budget Officer recently observed:

I think there is more or less general agreement that the principles that dominated textbooks on taxation have proven to be durable. Economists continue to talk about fairness in terms of horizontal and vertical equity. The framework to examine fairness – ability to pay, who benefits, fairness between generations and jurisdictions, and process fairness – remain useful and largely intact (Page, 2020).

Certainly, Canada has a long and continuing tradition of deference to the tax principles of horizontal and vertical equity. Our Royal Commission on Taxation was explicitly guided by those principles, and recent policy reform proposals continue that tradition. For example, the

---

<sup>3</sup> "This old approach is no longer prominent in academic economics, however. After James Mirrlees (1971) introduced the incentive compatibility approach to the taxation problem, combined with a utilitarian social welfare function, what economists do about the taxation problem is ... invoke an objective defined in terms of a final distribution of utilities, and they try to determine what tax system would induce the best distribution after individual agents react to the incentives laid out by the tax formula" (Fleurbaey, 2018, p. 38).

Department of Justice recently asserted that that “the harmonization of tax legislation cannot ignore the principle of horizontal equity” (Mandeville, 2022) and the 2021 Report on Federal Tax Expenditures identifies the realization of both horizontal and vertical equity as important policy objectives. Or consider David Duff’s influential overview of Canadian income tax law, which reflects an unqualified adoption of the principles Fleurbaey suggests are dusty relics:

“In evaluating the technical tax system, tax policy analysts generally turn to three criteria: equity, neutrality, and administrability. **The concept of equity has two components: horizontal and vertical equity.** Horizontal equity requires that individuals who are similarly situated be treated the same. The difficulty arises in determining when two individuals are similarly situated for tax purposes. The appropriate reference point, given that the income tax is designed to be levied based on taxpayers’ income, is income. Vertical equity requires that taxpayers with differing incomes be taxed differently. This aspect of the equity criterion draws on the concept of ability to pay. A taxpayer who receives greater income has a greater ability to pay than a taxpayer with lesser income, and therefore the income tax system ought to require the higher income taxpayer to pay more tax.” (Duff et al. *Canadian Income Tax Law*, 6th ed (Toronto: Lexis, 2018, emphasis added)).

Yet if Murphy and Nagel are right about their central claims, this is a rather striking state of affairs. Taxation is arguably our most important mechanism for realizing a conception of distributive justice, yet in at least some important cases the principles guiding that mechanism are fundamentally flawed. Simply put, if they are right, then there is a deep incoherence in conceiving of tax justice in terms of distributing a tax burden. As I mentioned, however, most people do not think they are right. I started this project similarly skeptical, and so studied the various objections in the literature to understand why. Surprisingly (at least to me) I have come to believe that these objections misunderstand the central thesis of *The Myth of Ownership*, and moreover that that thesis is more robust than is commonly recognized.

Part of this project is therefore an attempt to defend what I think is Murphy and Nagel’s most important conclusion—the incoherence of entitlement to the pretax baseline—against its critics. The short version is that critics are mostly concerned to reject Murphy and Nagel’s view

that property is merely conventional, and those critics seem to assume that such rejection is sufficient to rescue at least some entitlement to the pretax baseline. This strategy is understandable, given that Murphy and Nagel themselves appear to logically tie their criticism of the pretax baseline to conventionalism about property. Yet my own view is that these are two independent theses, and there is good reason to reject entitlement to the pretax baseline as incoherent *whatever* we think about property. Admittedly, this conclusion places me among the minority view, along with Lawrence Zelenak, who (I think rightly) describes Murphy and Nagel as “Cassandra-like, they know they are right, but that few will believe them” (2003, p. 2263).

A second part of this project is my argument that a remarkably similar issue plagues concepts fundamental to the other pillar of normative tax analysis—efficiency. I follow Murphy and Nagel in thinking that core concepts of tax fairness such as horizontal and vertical equity depend on a spurious distribution of entitlements as a baseline for making comparative judgments. But I also extend this argument, observing that core concepts of tax efficiency such as deadweight loss and Pareto optimality exhibit a remarkably similar pattern. They are necessarily comparative, and all too easily rely on spurious baselines for making comparisons. For instance, Pareto efficiency is definitionally comparative and so necessitates a comparative baseline. Where the status quo is arbitrarily relied on to provide that baseline, the Paretian analysis only applies an equitable-seeming formula (“some change makes everyone better off”) to a morally arbitrary initial distribution (“compared with the status quo”). A similar baseline problem is evident in the concept of deadweight loss, which wrongly implies that taxes inherently reduce economic efficiency. Here the comparative baseline is an ideal market with no taxes, but that is a misguided counterfactual if the objects of exchange are legally defined and thus the concept of a transaction is meaningless without a state and its laws.

A third part of this project emerged from my work on a particular Canadian tax credit (the Disability Tax Credit). Through this work I was exposed to the practice and justification of tax expenditures, which seem to me a perfect illustration of how problematic baseline thinking is entrenched in institutional practice, even when that practice is widely seen as exemplary and progressive. Moreover, I want to illustrate that the way politicians and everyday people discuss taxation and tax expenditures is utterly fettered by the view that private markets create moral entitlements with which the state interferes after the fact. It is difficult even to imagine the character of an alternative discourse. Nevertheless, I offer a suggestion to that end.

To my mind, salvaging the conceptual argument against entitlement to some pretax baseline opens the door to a different way of thinking about tax justice. The orthodox view of taxation as a process of state interference with an existing distribution of rights should be abandoned, along with the misguided project of identifying some “equitable-seeming function” and then applying it to some arbitrarily described distribution of property rights. Yet its recession makes space for a better way of thinking about tax justice. Taxation can be conceived as a tool of social architecture—a powerful means of constructing resource entitlements consistent with what justice requires. Indeed, it may be our most important tool for realizing distributive justice. With this alternative conception in mind, perhaps we can better realize the promise of taxation.

## **Project Overview**

The purpose of this project is to articulate and defend the view that both fairness and efficiency as they are understood in a great deal of tax analysis are plagued by a similar kind of comparative conceptual error, which I characterize as a baseline problem.

The first and second chapters set out the relevant background literature and main concepts necessary to make the case that tax analysis suffers from a baseline problem. The following three chapters then elaborate that problem, both in theory and practice. Chapter three turns to concepts of tax fairness, chapter four to concepts of tax efficiency, and chapter five to the practical contexts of public policy and public discourse.

The first chapter argues that property is a social relation among people. However, that relation is not easily characterized, and the standard triadic view of “A owns B against C” wrongly suggests a sharp distinction between the “things” and the “relations” of property. In addition, that standard view is parasitic on a remarkably indeterminate conception of ownership. Indeed, the legal practice of property largely ignores the concept of an owner altogether, and today property is “everywhere embodied in law” (Harris, 1996/2002, p. 3). Attention to this practice reveals how concepts like “pretax income” do not reflect a hallowed market baseline with which taxation interferes. Instead, they are constructed by the legal system as preliminary steps in a calculation of legal entitlement.

The second chapter considers several prominent attempts to theorize taxation as a normative inquiry. I divide these attempts into two broad categories. The first category takes seriously the idea that preliminary legal descriptions of pretax holdings represent fixed moral points against which the interference of taxation can be assessed. This tradition reaches from Ancient Greece to the modern public finance orthodoxy. It offers normative principles for

evaluating the impact of taxation on pretax holdings. The two most influential principles are that taxes should be responsive to ability to pay, or to the benefits people receive from the social system. Both principles assume that legal descriptions of pretax holdings provide a ready baseline against which to measure the identified dimensions of moral relevance. However, if that baseline is not morally significant, then both principles are only “applying some equitable-seeming function to a morally arbitrary initial distribution (Murphy & Nagel, 2002, p. 30).” In contrast to this general approach is a second way of thinking about taxation which I characterize as the social architecture view. This view takes seriously the role of the law in constructing descriptions of pretax holdings and maintains that such descriptions cannot provide fixed moral points for assessing tax justice. The general justice of a distribution like pretax income could only hold if tax justice was already secured, since taxation is an essential precondition of that distribution. But that, it has been argued, represents a kind of conceptual impossibility. In addition, the social architecture view denies the validity of tax justice as an independent domain, insisting on a holistic assessment of distributive justice. This view shifts the focus of justice from the burdens of taxation to the outcomes of the overall social system, including taxation.

The third chapter focuses on the “kind of conceptual impossibility” raised in chapter two, with the goal of reinforcing the diagnosis of a baseline problem in some prominent conceptions of tax fairness. Here I begin by introducing Murphy and Nagel’s infamous conceptual argument against the view that any pretax distribution of property rights morally constrains taxation. I survey various influential objections to their argument but conclude that those objections respond to a regrettable conflation in *The Myth of Ownership* between the conventionality of property and the incoherence of entitlement to a pretax baseline. My view is that these points are independent of each other, and that the conceptual argument does not rely on a deep conventionalism about

property rights. I propose an alternative interpretation of conventionalism limited to legal descriptions. This separation rescues the conceptual argument from its critics and poses it once again as a powerful critique of the orthodox principles of tax fairness.

The fourth chapter extends the baseline problem analysis to efficiency, the other pillar of normative tax analysis. I claim that two core concepts--Pareto efficiency and deadweight loss--are subject to a similar criticism as the one which Murphy and Nagel apply to the orthodox principles of tax fairness. That is, both are normatively attractive functions (often) applied to morally arbitrary initial distributions.

The Pareto principle recommends resource allocations that at least some people prefer and none disprefer. The principle has wide support, since it exhibits a kind of neutrality: it gives people more of what they want, by their own lights, without making anyone worse off. It is hard to imagine much resistance to Pareto-improving changes. However, I argue that this standard presentation is misleading. It masks a key consideration: an allocation people prefer to *what*? The baseline for comparison is essential, and the status quo is often smuggled in as the relevant baseline. But then the Pareto principle only recommends Pareto-improvements as compared with the status quo. Crucially, it does not recommend Pareto-improvements over the status quo over allocations that are Pareto-incomparable with the status quo. I have no argument against preferring Pareto-superior allocations to their Pareto-inferior comparators. But this rests on entirely different normative grounds from the claim that Pareto-improvements over the status quo should be preferred to allocations that are Pareto-incomparable to the status quo.

Deadweight loss is another key concept in analyses of tax efficiency, and it too has a baseline problem. In the tax context, deadweight loss refers to unrealized economic gains caused by taxation distorting market behaviour. Fewer taxes means more mutually beneficial

transactions, and so greater realization of the potential market surplus. At the limit, an untaxed market rests at a Pareto-optimal equilibrium, and the costs of taxation can be measured as deviations from that ideal market. While the untaxed market might be a useful fiction in economics, it is a mistake to apply the concept too readily to real world contexts. It implies that taxation creates an economic loss as compared with an ideal untaxed market, but that is simply false. As chapter one explains, the law (including taxation) is essential in providing the legal property structure that makes modern markets possible. The point also extends to contract, tort, and other elements of the law without which nothing like a modern market would exist (see, e.g., Pistor, 2019).

The final chapter takes up two case studies to illuminate problematic baseline thinking in practical contexts. The first turns to tax policy, where the celebrated practice of reporting tax expenditures provides a ready example of a baseline problem. The practice began with the observation that government can use either taxes or spending to make identical changes to economic outcomes. In other words, it aimed to expose the practice of spending by not taxing. But spending and not taxing are only the same relative to a baseline of entitlement, which led to the criticism that tax expenditure analysis assumes government is entitled to the “last penny.” Tax expenditure enthusiasts reject that criticism because they claim government is only entitled to some baseline description of taxation. This baseline was originally called the normative tax, but modern versions tend to deemphasize its normative content. Yet, notwithstanding that tendency, baseline descriptions of the tax system always involve evaluative judgments about what constitutes the “core” or “structural” functions of the tax system in contrast with “special” or “preferential” provisions. To defend this assessment, I consider two frameworks, one from Canada and the other from the United States, and argue that both are unavoidably and profoundly

normative. I therefore suggest that, instead of “silence and halftruths” (Bittker, 1969) about the normative content of tax expenditure analysis, theorists consider candid normativity. This would involve identifying, articulating, and justifying the normative content of the baseline necessary to measure tax expenditures. I argue this practice could remain neutral and principled as tax expenditure analysts require.

The second case study considers Canadian public discourse about tax justice. This goes well beyond public policy, and can be found in the halls of parliament, in newspapers, in conversations at workplaces, and so on. I observe that a single animating ideal dominates (of course, not exclusively) this discourse: paying one’s fair share. Support for this ideal is so widespread that I suggest it may be vacuous. I acknowledge that this criticism probably goes too far, since the same might be said of justice or other indispensable but contested concepts of which there are many conceptions. Therefore, I suggest reinterpreting our conception of the fair share as *receiving* a fair share rather than paying one. This reinterpretation better accords with the theoretical observations made in this project, most crucially that the legal description of pretax holdings does not provide a morally relevant baseline from which people could pay anything at all. Instead, what matters is what people *have*, after tax. That is the correct standard for distributive justice, but it stands in stark contrast to the modern discourse about taxation and fair shares.

## Chapter 1: Mine and Thine

### 1.1 Chapter 1 Introduction

There are two main reasons why I think evaluating the normative dimensions of taxation first requires some direct consideration of the concepts of property and ownership. The first is the general observation that both concepts are deeply contested and yet taxation cannot be studied without their constant reference. Glossing over the many conceptions resulting from this contestation can generate confusion about the relationship between property and social institutions and therefore has downstream implications for analyses of tax justice. Indeed, this is evident in the majority of arguments about tax justice which rely on these concepts without even roughly clarifying their intended meaning. The second reason is that canvassing certain aspects of several influential conceptions of property and ownership supplies the conceptual background for my central argument in chapter three, which relies on distinguishing “pretax” income as a legal property construct from other conceptions of property.

I emphasize only “certain aspects” of these conceptions because I cannot provide a comprehensive overview of property or ownership. Instead, I want to chart a path through some relevant parts of the literature that I think are especially important for my arguments about normative tax analysis. To that end, I aim to (somewhat artificially) separate the distributive from the conceptual aspects of property, where the former is concerned with “*what* is it that you own” and the latter with “what it *means* to own it” (Attas, 2006, p. 125). Though the two questions are never truly separable, engaging with the distributive question seems futile in the absence of reasonably developed conceptions of property and ownership. That is, if the conceptual analysis of property is in some ways prior to questions of distribution, then “[d]istributive justice must concern itself not just with ‘who has what’ but also with the nature of

‘having’” (Christman, 1994, p. 4). Understanding the ‘nature of having’ or ‘what it means to own’ requires some account of the contents or *structure* of ownership and, relatedly, of property.

To that end, I begin by introducing Thomas Merrill’s characterization of property as a social strategy for assigning control over resources. I then characterize conceptions of property as existing on a kind of spectrum of moral depth. At one extreme is the view that property rights are by their very nature inconsistent with taxation; for example, Attas (2006, p. 125) argues that “the simple fact [is] that a property right cannot survive its expropriation.” At the other extreme is the view that property rights are merely social conventions; for example, Murphy and Nagel write that property rights are “created by the tax system” (2002, p. 175), “purely conventional” (2002, p. 175), and that “individual citizens don’t own anything except through laws” (2002, p. 176). Somewhere in between lies the orthodox view that property rights are moral entitlements, shaped by social institutions, and present an important but defeasible constraint on taxation.

I also consider how property relates individuals to resources. Here I accept the largely consensus view that property rights are among persons rather than between persons and things. But I note that the nature of that relationship among persons has been disputed for centuries. Today, the commonsense view is that someone (an owner) has some unified rights (ownership) to something (property) against everyone else. Yet this is in contrast with over a century of theorists, often lawyers, insisting that the *practice* of property lacks any unified core of rights. Rather than a settled package of rights following from the identification of an owner, the title “owner” emerges as a legal honorific bestowed *after* the assignment of some (often idiosyncratic) bundle of rights. According to this view, the concept of property must be fully disintegrated, and its various elements only reassembled on a case-by-case basis. However, against disintegration, several theorists in recent decades have argued that property implies a

unified core of rights after all—even in the law. Thus, even academic conceptions of property have drifted from unity to disintegration to perhaps a kind of unity once again.

My objective here is not to side with a particular “unified core” or “disintegration” view of property. Instead, it is to reveal and emphasize the sense in which certain notions of “pretax” entitlement refer to constructs that are profoundly legal and so take an antecedent tax structure for granted. This in turn grounds a serious skepticism of anything like the Lockean view as a paradigm for thinking about property, ownership, and taxation. The widely held view I want to reject looks something like this:

Someone has an extensive set of entitlements to a bushel of apples after picking them from a tree. The person can consume them, give them away, destroy them, and so on. These are moral entitlements, perhaps because the person cultivated the orchard, needs the apples for nutrition, entered into a certain agreement with someone who has the relevant entitlements, or whatever. Legal rules are simply social constructs that track and protect these moral rights.

Given such a view, one might conclude that entitlement in a modern society to pretax income is analogous to having moral rights to the bushel of apples. That is, in both cases people have moral rights to certain specific resources (apples, income) and the protections of the law are simply layered on top of that moral structure as best they can be. Taxation might interfere with that moral structure by requiring one person to give some apples to another or taking someone’s income and giving it to another. This interference may be justified but it is interference with some existing moral entitlements nonetheless. This is more or less the commonsense view of taxation, and as shown in chapter two it is inherent in many of the concepts used in assessing tax fairness.

Against this view, my characterization of pretax income as an obviously legal construct implies that there is a significant disanalogy between pretax bushels of apples and pretax income. Someone could be morally entitled to a bushel of apples in the absence of the Canadian state, but

the same cannot be said for entitlement to \$50,000 Canadian. That is, pretax income is a legally derived calculation, constructed in accordance with the relevant tax laws and regulations (for example, the definition of income), used as a baseline for calculating other legal entitlements. Moreover, and unlike the apples, pretax income does not pick out some distribution of resources that people writ large even *could* be entitled to *against* taxation. This is because the “pretax” figure is not, in fact, pretax at all: it depends on the state (and so taxation) for its character and very existence. As such, any entitlement to it would depend on antecedent existing and justified taxation practices.

In summary, ostensibly “pretax” constructs like pretax income are not antecedent to taxation at all. They are profoundly legal, and there is no reason to assume that they reflect moral entitlements that constrain taxation. Since such constructs are outputs of a larger system of which taxation is an essential part, taxation must be justified before the resource allocations they pick out (outputs of that very system) could carry moral weight, even in principle. The debate between unified core and disintegration theorists is useful because it clearly demonstrates the role of the law in creating rather than tracking some description of entitlements. As a result, arguments about tax justice err when they appeal to the various concepts of pretax holdings as constraints on what a justified tax system might look like. Such arguments conflate preliminary legal descriptions with moral property rights.

## **1.2 The Property Strategy**

For better or worse, we live in a society that broadly endorses “the property strategy” (Merrill, 2012, p. 2061). Every society adopts some combination of strategies for organizing control over resources. But the property strategy is distinctive for its emphasis on decentralized

control, as compared with, for example, bureaucratic governance or group consensus (Merrill, 2012). At its most abstract, it structures control according to a triadic relation: “‘A owns B against C,’ where C represents all other individuals” (Cairns, 1935, p. 59). Each element of this relation stands in need of significant clarification, but the core idea is that the property strategy assigns a domain of control or unique prerogative over a specific, discrete resource to a natural or juridical person. The content of ‘owns’ in the triadic relation (if any) has of course been contested for centuries; indeed, Jeremy Waldron submits that ownership may be an *essentially* contested idea (1990, p. 54).<sup>4</sup> For now, I suggest that the following statements capture the general character of the property strategy:

[O]wners are in a special position to set the agenda for a resource (Katz, 2008, p. 278).

The owner of a resource is simply the individual whose determination as to the use of the resource is taken as final (Waldron, 1990, p. 39).

[Property is] a method of assigning to particular individuals the ‘authority’ to select, for specific goods, any use from a nonprohibited class of uses (Alchian, 1977, p. 130).

[T]he owner has a wide—indeed, open-ended—set of options about what to do with the discrete resource (Merrill, 2012, p. 2069).

The central core of the notion of a property right in  $X$  ... is the right to determine what shall be done with  $X$ , the right to choose which of the constrained set of options concerning  $X$  shall be realized or attempted (Nozick, 1974, p. 171).

Again, none of these statements is sufficient to fully characterize, let alone defend, a conception of property or ownership. However, a precise definition may be unnecessary, if Gerald Gaus is right that “we typically have a far better grasp of the use of an idea than we have of the

---

<sup>4</sup> His reference is to W.B. Gallie (1956, p. 169): “. . . I shall try to show that there are disputes, centred on the concepts which I have just mentioned, which are perfectly genuine: which, although not resolvable by argument of any kind, are nevertheless sustained by perfectly respectable arguments and evidence. This is what I mean by saying that there are concepts which are essentially contested, concepts the proper use of which inevitably involves endless disputes about their proper uses on the part of their users.”

philosophical analysis of it” (2010, p. 242). For now, I only submit that these statements pick out an intuitively familiar fact about the relations among people with respect to resources. In our society, there are certain resources, normally those that are discrete, scarce, and valuable, over which specific individuals have a unique scope of control. According to the property strategy, if the umbrella is mine then (within the scope of permitted uses) I can do with it whatever I like. You may need it more, want it more, deserve it more, be able to make better use of it, or whatever. Regardless, it belongs to me.

Every known society follows the property strategy “to one degree or another” (Merrill, 2012, p. 2062), and there are a number of reasons why. Merrill lists various advantages, such as the value of local knowledge, incentives to maximize value, scalability, avoiding tragedies of the commons, providing a foundation for exchange, checks and balances, and facilitating the expression of personhood. In fact, advantages of the property strategy are identified even in ancient texts, such as Aristotle’s *Politics*. There he writes, “that which is common to the greatest number has the least care bestowed upon it” (II.III). According to Aristotle, individuals favour their own interests, while thinking “hardly at all of the common interest.”<sup>5</sup> Thus he concludes that assigning control over resources to specific individuals promotes progress because “everyone will be attending to his own business” (II.V). Of course, the property strategy also has its disadvantages. Plato associated it with moral corruption in his *Republic*, and modern theorists like Merrill (2012) note that it brings new externalities, risks of monopoly, inequality (which critically undermines its advantages), and so on.

Presentation of the property strategy in terms of its advantages and disadvantages casts both its scope and content as flexible and at least partially indeterminate. A society might use it

---

<sup>5</sup> This point was later made famous by Garrett Hardin’s (1968) *The Tragedy of the Commons*, though without reference to Aristotle.

to organize control over personal property but not productive property, or over chattels but not land. It might assign control over B to A, or it might assign control over B to C, whichever best suits the ends the strategy is intended to realize. Similarly, the content of ‘assigning control’ seems partly up for grabs. A society might assign some combination of rights and liberties to possess, exclude, or alienate some resource. Or it might assign a quite different control structure, all depending on its aims. Accordingly, taxation seems to fit comfortably with the property strategy, at least in principle. Naturally, there may be tensions in practice. For example, if the property strategy is used because it creates incentives to maximize value, but some proposed tax undermines that benefit then, to that degree, taxation may be in tension with the property strategy. But there is no principled conflict; it is just an empirical question of how best to design and balance both taxation and the specific details involved with implementing the property strategy in a particular social context.

In contrast, many people do not think that society is free to decide on instrumental grounds when to select the property strategy from a menu of options, and nor do they think its form is as malleable as the preceding suggests. Instead, they think both the scope and even the form of the property strategy are dictated, or at least seriously constrained, as a matter of right. Property is an institution of deep moral rights on this view, and those rights limit what a system of taxation can rightfully do. On the most rigid interpretation, taxation involves taking what people have a right to, and so “involuntary taxation is inconsistent with respect for property rights” (Attas, 2006, p. 119). This alternative interpretation suggests that there is, after all, a principled conflict between property rights and taxation.

### 1.3 Rights or Conventions

The rough characterization introduced so far reflects a longstanding disagreement about the nature of property, and what might be called the moral depth of property rights. On one side are natural rights theorists, who argue that property rights are morally deep, fundamental even. On the other side are conventionalists, who argue that property rights are more shallow: useful conventions in service of other values. Despite this disagreement, we should note that both views accept that property rights are significant, perhaps even essentially so. As Bret Boyce explains, “historically the right to private property has been regarded as the central paradigm for rights in general, and the essential precondition for the creation of a private sphere of autonomy that forms the foundation of the pluralistic liberal order” (2007, p. 202). Suppose conventionalists believe that a private sphere of autonomy is valuable for its own sake. If they accept that private property is an essential precondition for realizing that private sphere of autonomy, then property rights are morally essential even for conventionalism.

To better appreciate the scope of disagreement between natural rights theories and conventionalism, consider a brief description of two prominent figures, each representing one of side of the dispute: John Locke (1690) and David Hume (1739). Most relevant here is their divergent understanding of the relationship between property rights and society. For Locke, property rights exist prior to society, and society must respect those rights. For Hume, there are no property rights prior to society. People may have possessions, but rights to those possessions are only created by society. A short way to summarize this aspect of their disagreement is that according to Locke, society preserves property rights. According to Hume, it creates them.

The best illustration of Locke’s view is his claim that “[t]he great and chief end, therefore, of men’s uniting into commonwealths, and putting themselves under government, is the

preservation of their property” (1690, IX). Note that his emphasis on *preservation* is consistent with his pre-social theory of property. If the principal end of creating society is the preservation of property then, presumably, there must already exist property to preserve. But then the pre-social existence of property itself calls for explanation, and Locke answers that call with his unilateral theory of appropriation, the labour theory of property:

Though the earth, and all inferior creatures, be common to all men, yet every man has a property in his own person: this no body has any right to but himself. The labour of his body, and the work of his hands, we may say, are properly his. Whatsoever then he removes out of the state that nature hath provided, and left it in, he hath mixed his labour with, and joined to it something that is his own, and thereby makes it his property (1690, V).

So for Locke, individuals outside society have more than mere possessions, they have property. Uniting in civil society extends the natural rights people have to themselves, their labour, and their possessions. These rights exist prior to society, and the chief end of governance is to preserve them.

At the other end of the spectrum lies Hume’s conventional story of property. Hume agrees with Locke that property is the main purpose of social organization. He describes property as the “chief advantage” and “the most necessary to establishment of human society” (1739, II.II.II). But he disagrees about the possibility of pre-social property. For Hume, there is no property in the “savage and solitary condition” (1739, III.II.II).<sup>6</sup> Instead, property “is nothing but those goods, whose constant possession is established by the laws of society; that is, by the laws of justice” (1739, II.II.II). Hume emphasizes that conventions of respecting constant possession do not arise from a promise or agreement but are rather the product of mutual recognition of advantage. He adds that if individuals were perfectly altruistic and resources were abundant there

---

<sup>6</sup> Likewise, Thomas Hobbes (1651, XIII): “It is consequent also to the [pre-social] condition, that there be no Propriety, no Dominion, no Mine and Thine distinct” and (1651, XV) “with the constitution of a civil power . . . propriety begins.”

would be no possibility of mutual advantage and so no property. But since neither is true, property is a socially useful response to at least moderate scarcity and at least partial egoism. What about original acquisition? Again, the “misery of the condition” that precedes society motivates people to act quickly in restraining themselves, and so they simply accept first possession as the proper basis for justifying successive transfers (1739, II.II.III). For Hume, property is “a convention entered into by all the members of the society to bestow stability on the possession of . . . external goods” (1739, II.II.II).

Hume understands himself as offering a certain kind of justification, despite his account possibly seeming like a mere description. This is because if, as he argues, private property is “necessary for *any* non-savage, non-chaotic form of social life, then there is no obvious need for any further justification of that system” (Panichas, 1983). The institution of property is thereby justified as a necessary condition for human social life. There remains the need for a justification of particular property conventions (specific rights, duties, liberties, and so on). Here Hume appeals to various utilitarian considerations, at least insofar as they do not require actions that might undermine the prior justification for the system itself (Panichas, 1983).<sup>7</sup>

The high theoretical stakes of the debate between Lockean-natural-rights-theorists and Humean-conventionalists is also thought to have profound political implications. As James Harris puts it, “if there are any natural property rights, that would radically affect the moral background against which property institutions operate” (2002, p. 228). The combination of theoretical interest and sweeping political implications has placed property in company with a

---

<sup>7</sup> Hume’s relationship to utilitarianism remains a significant source of controversy (e.g., Reichlin, 2016). Hume also denies the direct role of utilitarian thinking, as he argues that such a rule would attract the very uncertainty and disagreement that the institution of property is adopted to resolve (1739, Book III.II.II). The main point here is only that Hume appeals to various normative considerations that he sees as independent of the particular convention under consideration.

number of concepts which, taken together, construct the conceptual terrain for various political proxy wars over the distribution of resources (see, e.g., MacLean, 2017). Daniel Russell makes this observation about the related idea of self-ownership: “for decades ostensible debates over self-ownership have in fact been a front. The real debate beneath the surface has been an ideological debate over government-enforced redistribution of wealth, income, and property” (2018, p. 26). Even within the libertarian tradition, debates about world-ownership and initial acquisition can be interpreted as something of a ‘front’ since, according to some, “the real dispute between left and right is over taxation” (Fried, 2004, p. 80).<sup>8</sup> Arguments about the concept of property may be, like those about self-ownership and acquisition, inseparable from their political and distributive implications.

One prominent example of the political implications of the Locke-Hume divide comes from American constitutional history (Boyce, 2007). Participants in the 1787 Constitutional Convention made a number of statements which echo the view shared by both Hume and Locke that property is the primary purpose of government. They therefore emphasized the centrality of protecting property to the constitutional project. For example:

Alexander Hamilton: “One great Objt. of Govt. is personal protection and the security of Property”

Gouverneur Morris: “Life & liberty were generally said to be of more value, than property. An accurate view of the matter would nevertheless prove that property was the main object of Society.”

John Rutledge: “[p]roperty was certainly the principal object of Society”

---

<sup>8</sup> Debates about taxation are sometimes fronts for yet other debates. For example, Christopher Petrella (2017) explains how the relationship between taxation and slavery threatened the Philadelphia Convention. Participants were friendly to the ‘no taxation without representation’ maxim yet this forced them to make a choice—not even a month after declaring that ‘all men are created equal.’ If slaves were property, then slave ‘owners’ would face significant taxation. But if slaves were not property, then the legitimacy of slavery was called directly into question. So it was that Thomas Lynch remarked, “If it is debated, whether their slaves are their Property there is an End of this Confederation” (see also Einhorn, 2006).

Pierce Butler: “Government . . . was instituted principally for the protection of property”

Charles Pinckney: government was “instituted for the protection of property”

But did these men see the state as a tool for *preserving* property (as Locke claims) or did they see it as creating new conventions of property that could not exist without the state (as Hume claims)? The evidence suggests that the founders fell squarely on the Lockean side of the divide. In his well-known analysis, Charles Beard concluded that the U.S. Constitution is “essentially an economic document based on the concept that the fundamental private rights of property are anterior to government and morally beyond the reach of popular majorities” (1913, p. 324; see also McGuire, 2003). The key idea—at the very core of the U.S. constitution—is that property is *anterior* to civil society and presents a serious constraint on taxation. The merit of that notable idea is the focus of my third chapter.

In summary, the Lockean natural rights tradition is thought to establish property rights anterior to government, while Humean conventionalism maintains that property rights are created by the state. It is normally assumed that these views have very different implications for tax fairness. For Locke, natural property rights pose a constraint on what government (and everyone else) can do. Accordingly, the role of the government is to extend and protect those rights, and so those rights present a serious justificatory hurdle for taxation. For Hume, the moral status of property rights is always derivative of deeper values. In general, property may or may not be instrumental in realizing those values. Modern interpretations deviate from this stylized presentation of the Locke-Hume divide, but there is a general consensus that the moral status of property rights—especially whether or not they are conventional—is a question of major significance for tax fairness. On this point, a Lockean view of property continues to exert a powerful force on the common sense of taxation.

## 1.4 Conceptions of Property and Ownership

Does property exist anterior to civil society? Does ownership present a moral constraint on taxation? Are property rights derivative, conventional, and malleable, or fundamental, natural, and rigid? Trying to answer these and other similar questions relies on some background assumptions about what we mean by “property” and “ownership.” Sometimes usage refers to legal conventions, other times to what an author believes are deep moral rights. Naturally, as Harris observed above, these divergent views are thought to lead to similarly divergent conclusions about tax justice. What I argue in chapter three, however, is that neither view can ground entitlement to ostensibly pretax legal descriptions such as pretax income. That argument depends on recognizing how thoroughly our conception of entitlement is constructed by legal rules, *including taxation*, which themselves stand in need of justification. It depends ultimately on what is meant by “my income, before tax.” The notion of “my” (or *mine*) evokes a familiar sense of entitlement (ownership) over some thing (income). But this is deceptively simple, as for centuries people have struggled to clarify what exactly it means for some *thing* to be *mine*.

### 1.4.1 Foundations

The particularities of contemporary conceptions of property and ownership are descendants of a foundational distinction between *meum* and *tuum* (mine and thine). In ordinary language the notion of *mine* is expressed in the “obvious distinction . . . between having a thing and being entitled to have it” (Nicholas, 1962/2008, p. 107). Or, as Rousseau put it, having a thing implies only *possession*, while it being mine implies an entitlement relation which is called *property* (1762, Book 1.8). Practices of distinguishing mine from thine, and so of possession and entitlement, are clearly evident by the Neolithic period, and probably are as ancient as human

organization itself (see, e.g., Ellickson, 2013; Krier, 2009; Bowles, 2006). Indeed, Merrill argues that the universally practiced convention of gift giving implies a shared understanding that specific things in some sense belong to particular people (2012, p. 2073). The influence of this elemental distinction can hardly be overstated—perhaps no other idea so formidably shapes social organization. Honoré thought that a people without ‘mine and thine’ would be alien to human society: they would “live in a world that is not our world” (1961, p. 370). Blackstone (1753) famously asserted that “[t]here is nothing which so generally strikes the imagination, and engages the affections of mankind, as the right of property” and when Thomas Hobbes turned his mind to the theory of government, his first subject was how “any man should call any thing rather his Owne, than another man's” (1651). Rousseau cynically observed that “[t]he first man who, having enclosed a piece of ground, bethought himself of saying This is mine, and found people simple enough to believe him, was the real founder of civil society” (1755). As I previously noted, even the diametrically opposed understandings of Locke and Hume nevertheless converged on property as the first virtue of social organization, and so too did their political descendants.

Yet for all its age, familiarity, and significance, there is no widely accepted definition of property or ownership; there is no consensus on how to cross the “unbridgeable logical gap between the fact of possession and the right to possession” (Epstein, 1979, p. 1240). There is an inexhaustible theoretical literature, and there is also a social practice of property evident both in law and in everyday life. These practices both inform and are to some extent animated by what might be called the common sense of property. Attention to these practices is essential to reducing the risk of adopting a merely stipulative account. For example, it might be tempting to follow Bentham’s oversimplification of ‘all property as law.’ But, as I later detail, this would

leave unexplained his exclusion of all non-legal practices that exhibit what are today accepted as core features of property. A better approach is the “rational reconstruction” (Habermas, 1979) of property and ownership, beginning with experience and common sense (James, 1907, p. 24).

This is not because common sense is definitive, but rather it reflects a modest theoretical conservatism, as David Lewis explains:

Common sense has no absolute authority in philosophy. It's not that the folk know in their blood what the highfalutin' philosophers may forget. And it's not that common sense speaks with the voice of some infallible faculty of 'intuition.' It's just that theoretical conservatism is the only sensible policy for theorists of limited powers, who are duly modest about what they could accomplish after a fresh start. Part of this conservatism is reluctance to accept theories that fly in the face of common sense, but it's a matter of balance and judgment (1986, p. 134).

There is of course no universally 'common' sense of property and ownership since different cultures in different times have genuinely distinct means of organizing themselves with respect to resources. Still, I think it is reasonable to say, as Thomas Grey does, that today there is a commonsense understanding of property as “*things that are owned by persons*” (1980, p. 69).

### **1.4.2 Relations with Things**

The commonsense understanding identified by Grey provides a useful starting point, but it needs interpretation. One interpretation is that ownership is a relation between a person and some things. Indeed, people ordinarily speak of certain things as property: they say that umbrella or that house is their property. This suggests a dyadic structure that connects persons to things: A owns B. The implication is that, if A owns B, then B is property. Intimations of this dyadic structure interpretation are also found in some theoretical accounts of property. For example, David Hume writes that property is “a relation [between] a person and an object” (1739, II.I.X) and that property “consists in the relations of objects to intelligent and rational beings” (1739,

III.II.VI). The dyadic structure implies that an isolated person can have property on a deserted island; after all, the island has a person and things. Accordingly, Hume accepts that a person landing on the shores of a deserted and uncultivated island can rightfully acquire property in the entire island.<sup>9</sup> Another dyadic account comes from Charles Toullier who describes property as “a *moral quality* inherent in a thing; *an actual bond* which fastens it to the proprietor, and which cannot be broken save by his act” (Proudhon, 1840). A final example comes from Thomas Piketty, who describes property as “power over things” (2020, p. 5).<sup>10</sup>

One problem with interpreting ownership as a dyadic relation between a person and some things is that property is a moral (or legal) relation of entitlement. Such relations are necessarily among people, not strictly between people and things. As Jan Narveson (2010, p. 113) explains, ownership “is *not* a relationship between you and the thing owned, simpliciter. This seems not to have been appreciated by some writers, and one can only suppose that the reason is that from the point of view of the owner, it might look that way. But to think that is to confuse possessions with property.” A person can possess a thing, but the property relation is not merely a factual description of how that person relates to that thing. Justice Markby observed in *Elements of Law* that “if we attempt to translate the phrase [*in rem*] literally, and get it into our heads that *a thing*, because rights exist *in respect of it*, becomes a sort of *juristical person*, and *liable to duties*, we shall get into *endless confusion*” (Hohfeld, 1917, p. 76). That is, if we conceive of property in terms of rights, it seems clear that those rights must be against entities that can bear obligations, which excludes things. James Penner points out that “a concept like ‘possession’, [...] concerns

---

<sup>9</sup> Throughout his account Hume blurs dyadic and triadic structures since he also describes property as a relation that “permits [the person], but forbids any other, the free use and possession of [the object], without violating the laws of justice and moral equity” (1739, II.I.X).

<sup>10</sup> This was also the original understanding of rights *in rem* (literally rights against a *thing*). But all rights are against people, so that interpretation made little sense. Now *in rem* is usually understood to mean ‘against a large and indefinite class of people’ or, more generally, ‘the world.’ However, I also think Piketty could be interpreted in a non-dyadic way here if “power” refers to a relation among people.

the relation of a person to a thing, not a person's relations to others" (2020, p. 7). Kant also challenged the dyadic view on the same grounds in his *Metaphysics of Morals* (Ak. 6:261):

. . . it is clear that someone who was all alone on the earth could really neither have nor acquire any external thing as his own, since there is no relation whatever of obligation between him, as a person, and any other external object, as a thing. Hence, speaking strictly and literally, there is also no (direct) right to a thing. What is called a right to a thing is only that right someone has against a person who is in possession of it in common with all others (in the civil condition).

Kant considers the possibility that a property right is a direct relation to a thing, but he criticizes as "obscure" and "absurd" the idea of rights and obligations between persons and things (Ak. 6:260). This is because a right to possess an umbrella is really a relation among people with respect to that umbrella.<sup>11</sup> Kant therefore rejects the conception of a property right "as if it were a *guardian spirit* accompanying the thing" (Ak. 6:260). Instead, he plausibly argues that a right to a thing is a right against others with respect to that thing.

### 1.4.3 Relations with Persons

The implausibility of the person-thing dyadic account of property is now generally accepted. Property is a social relation, and there can be no property on a desert island; as Demsetz puts it, "[i]n the world of Robinson Crusoe property rights play no role" (1967, p. 347). Legal practice has long recognized the same. As early as 1934 Hans Kelsen, for example, "could express incredulity that the naïve view 'continues to be maintained' – 'despite the repeated objection that the legal power of a person over a thing consists in nothing else than a specific relationship to other subjects'" (Sage, 2018, p. 116). Contemporary theorists therefore tend to follow Cairns in describing property as a triadic relation: A owns B against C, where C is

---

<sup>11</sup> See also Alan Gibbard's (1976) assertion that "a *property right* is, roughly, a [moral] right which a person has with respect to a specific thing" (p. 77).

everyone other than A. This triadic structure captures a broad theoretical consensus that property entails a relation among agents with respect to resources. Here are some examples:

[O]wnership is a normative relation that cannot be identified with any physical property. To own something is not to be in physical possession or control of it, but to be *entitled* to have others forbear from interfering with it in specified ways (Ripstein, 2004, p. 27).

Any assertions about rights “to things,” as we are accustomed to putting it, are assertions about what we may do in relation to each other, with or to or about the things in question” (Narveson, 2010, p. 113).

[P]roperty does not exist on a desert island, there being neither the context nor the need (Lametti, 1998, p. 715).

Ownership is therefore not a relation between an owner and a thing, but between an owner and all other persons. Ownership comprises claims and liberties that define social relations; it is not a relation in which a Robinson Crusoe could stand. It is only because we live together that ownership is part of human life (2018, Russell, p. 27).

This looks contrary to the common-sense view that ‘B’ in the triadic relation is itself property. If property is the relation between A and C with respect to B, then B is not itself property. B is just a thing, the object of property. Some theorists simply accept this as a different use of the term property. They hold that the term ‘property’ is ambiguous between property-as-thing and property-as-triadic-relation, but that context normally indicates the intended usage.<sup>12</sup> One historical response to this ambiguity has been to legally distinguish between property as thing (tangible) and property as relation (intangible). It is worth pausing to consider why this failed, as it usefully illustrates the sense in which property rights always include both relations and things.

---

<sup>12</sup> See, for example, (Lametti, 1998, p. 665): “In both lay and legal discourse, the general term ‘property’ is used at once to refer to both the objects or subject-matter of property (i.e., property-as-thing or property-as-resource) as well as the ways in which it can be held or the factual or juridical relationship persons have to the objects of property (i.e., property-as-relationship).”

#### 1.4.4 The Gaian Inconsistency

The Institutes of Justinian<sup>13</sup> conceive of property as things owned, and ownership as a relation of *dominium* between the owner (*dominus*) and the thing (*res*). The category ‘thing’ includes the tangible (*res corporales*) and the intangible (*res incorporales*). Tangibles include “land, slaves, clothing, gold, silver, and others innumerable” (II.II).<sup>14</sup> Intangibles include “rights . . . such as inheritance, usufruct, and obligations” (II.II).<sup>15</sup> To see the problem in this scheme, imagine a Roman who buys one plot of land (tangible thing) and a usufruct in another (intangible right). The categories of the law imply that these cases are distinct, but this is misleading (Gretton, 2007). In the first case the Roman buyer does not acquire *the land* and in the second case some cluster of entitlements to the land. *Both* cases could be described in terms of entitlements or both in terms of the tangible objects governed by those entitlements. The purchase of ‘land’ is itself an acquisition of a cluster of entitlements with respect to the land. To describe it as qualitatively distinct from a usufruct is arbitrary, since ‘buying a plot of land’ is only elliptical for buying an entitlement that is “merely more extensive” than the usufruct (Nicholas, 1962/2008, p. 107).

Depending on its interpretation, the modern triadic structure risks importing the Gaian inconsistency. To see how this works, consider a typical case in which A purchases Blackacre.

The resulting relation can be input into the triadic form, A owns B against C:

*[A] owns [Blackacre] against [everyone else].*

---

<sup>13</sup> A 6<sup>th</sup> century codification of Roman Law, based largely on the institutes of the Roman jurist Gaius.

<sup>14</sup> Slaves are first discussed under the law of persons (Book I), but also included as corporeals under the law of things (Book II.II). Of interest is the qualification in Title I.III that “slavery is an institution of the law of nations, *against nature* subjecting one man to the dominion of another” (emphasis added). This is a specific instance of a more general trend in which slave societies recognized slavery as unnatural and wrong, even when slavery was widely practiced. For more, see Graeber (2011) 165 – 210.

<sup>15</sup> The term *usufruct* is a combination of *usus* (use) and *fructus* (fruits). These are (roughly) rights to possession and profits, with their combination (roughly) corresponding to the modern life interest.

In this case (A) holds ownership rights (whatever that entails) with respect to the thing (B) Blackacre, and those rights are held against (C) everyone else. But then consider a different case. In this case A does not purchase ‘Blackacre’ but instead purchases a usufruct in Blackacre. Now it appears that A has ownership rights with respect to a different thing, a usufruct in Blackacre:

*[A] owns [a usufruct in Blackacre] against [everyone else].*

But here we have stumbled directly into the Gaian inconsistency. In both cases, A purchased a package of rights with respect to Blackacre. Yet we arbitrarily described the purchased thing (B) as ‘Blackacre’ in the first case and as ‘rights in Blackacre’ in the second case. To resolve this inconsistency, we can proceed in either of two ways. The first resolution involves treating Blackacre as the object of the relation in both cases:

*[A] owns [Blackacre] against [everyone else].*

*[A] ‘holds’ usufruct in [Blackacre] against [everyone else].*

The problem with this resolution is evident by the placement of ‘holds’ in the reformulated second case. It seems like A stands in an important relation not only to Blackacre, but also to the usufruct in Blackacre. Namely, it seems like A purchased and now *owns* the usufruct. After all, in the normal case A can alienate or exchange the usufruct, set the agenda for it, and so on. But if A owns the usufruct, then consistency requires us to treat both packages of rights as the relevant object in the triadic relation:

*[A] owns [ownership in Blackacre] against [everyone else].*

*[A] owns [usufruct in Blackacre] against [everyone else].*

In the first case, A purchased and now owns Blackacre. But this formulation imagines ‘ownership in Blackacre’ as a package of rights now owned by A, which introduces an additional

layer of ownership. Thus, the following table summarizes the Roman law dilemma as applied to the triadic structure of property:

Table 1: The Gaian Inconsistency

<b>Consistency</b>	<b>(A)</b> Owner	<b>(owns)</b> Rights	<b>(B)</b> Thing
Inconsistent	A	ownership bundle	Blackacre
	A	ownership bundle	usufruct in Blackacre
(1) Consistent	A	ownership bundle	Blackacre
	A	usufruct bundle	Blackacre
(2) Consistent	A	ownership bundle	ownership in Blackacre
	A	ownership bundle	usufruct in Blackacre

If we resolve the inconsistency by (1), then we have to explain how A ‘holds’ the usufruct when it appears that A *owns* the usufruct. If we resolve it by (2), then we have to explain how A owns the package of rights called ownership, and the potential infinite regression (does A own the ownership rights over the ownership rights to Blackacre?).

Here is one more example to illustrate the point. An umbrella (a thing) will ordinarily be described as property. We can express that relation in terms of the triadic structure: [Henry] owns [the umbrella] against [everyone else]. Suppose Emma purchases ‘umbrella use rights’ from Henry. This creates a new triadic relation: [Emma] owns [umbrella use rights] against [everyone else]. This again illustrates the Gaian inconsistency. We have arbitrarily described [B] (the thing) as an umbrella in the first case and *rights* to an umbrella in the second case. Yet both

cases involve a thing (an umbrella) and a package of rights with respect to that thing, with Henry's merely being more extensive than Emma's. This inconsistency pushes us towards the first resolution, in which we describe the 'thing' in both cases as an umbrella and describe Henry as owning it while Emma has some other package of rights.

*[Henry] owns [the umbrella] against [everyone else].*

*[Emma] holds use rights to [the umbrella] against [everyone else].*

This has the virtue of consistency, but it fails to capture a crucial element of Emma's relation to the use rights. It claims she 'holds' them but that obscures the fact that they *belong* to her in a particular way—namely, she seems to own them. Those rights were the object of her purchase from Henry, not the umbrella itself. She purchased those rights, and now she seems to own them. This trouble pushes us towards the second resolution, in which we concede that the use rights are themselves the object of the relation.

*[Henry] owns [ownership of the umbrella] against [everyone else].*

*[Emma] owns [use rights to the umbrella] against [everyone else].*

The result is a kind of regression in which we must acknowledge that Henry owns the ownership rights to the umbrella. This may appear awkward, but layers of ownership do apparently track the common (and legal) practice of property. Consider that Emma's purchase of umbrella use rights *may* include higher order rights to, for example, sell those use rights, but they may not. Thus, there is a genuine question about the property rights a person has with respect to property rights themselves. It is an open question whether Emma also holds higher order rights regarding her use rights to the umbrella. It is true that she holds the right to possess the umbrella, but she may or may not hold the right to sell the right to possess the umbrella. Given this indeterminacy

in the notion of ‘owns’ it may be clearer to drop the concept of ownership from the relation and simply refer to a package of rights.

*[Henry] holds some rights to [the umbrella] against [everyone else].*

*[Emma] holds some rights to [the umbrella] against [everyone else].*

This suggestion will be considered more fully below, but first, what are the implications for the triadic formulation of property?

#### **1.4.5 Revisiting the Triadic Formulation**

Recall that the triadic formulation holds that “ownership is a normative relation with three terms: (1) you, the owner; (2) whatever you own; and (3) everybody else” (Narveson, 2010, p. 113). On the standard view, this entire relation is called ‘property.’ But this contributes to ambiguity, because (2) in the formulation is also ordinarily called ‘property.’ Moreover, the Gaian inconsistency illuminates why we cannot conceive of the ‘thing’ and the ‘relation’ as fundamentally distinct aspects of property: property always involves both relations and things. Can better language resolve the ambiguity? Leif Wenar argues that it can, and so we should “move away from the uncomfortable conflation of an object of rights with the rights themselves” (1997, p. 1944). His suggestion is that “private property is all those things over which private property rights are held” (1997, p. 1944). Accordingly, “far from all property rights being property, none are—they are just property rights. Property is what property rights are rights over” (Wenar, 1997, p. 1944). In general, then, his formulation is that A owns B against C, where A is the owner, B is the property, and C is everyone else. If A stands in this relation, then we can say that A has property rights in B. This proposal aims to remove the ambiguity from the triadic formulation without offending the commonsense view.

This ambiguity remains, however, as consideration of the Gaian inconsistency revealed how sometimes the objects of entitlement are themselves property rights. This is a problem for Wenar's view, since things can be *both* property rights and property in the sense he describes. The distinction between rights and their objects is not mutually exclusive. This is a key feature of the *summ*<sup>16</sup> in natural law, as described below, in 1.4.9: "In a word, as a man is said to have a right to his property, he may equally well be said to have a property in his rights" (Madison, 1792). For example, suppose I purchase the most extensive available package of rights over Blackacre. According to Wenar, those rights are property rights and Blackacre is property. Suppose I sell only the right to possess the land to Wenar. Now Wenar has the right to possess the land, but, like Emma and the umbrella, he may (or may not) also have the right to sell that right to possess. Notably the *right to possess* Blackacre is Wenar's property rather than Blackacre. Suppose Wenar does have the right to sell his right to possess Blackacre. Now he wishes to sell *the right to sell* the right to possess Blackacre. Henry buys the right to sell the right to possess—is (part of) Blackacre Henry's private property? Contrary to Wenar's claim ("far from all property rights being property, none are"), the objects of ownership and exchange (thus property on his account) are *also* property rights. We can adjust his proposal by allowing that in some cases property rights are property after all, but unfortunately this means the "uncomfortable conflation" between property rights and their objects has returned, at least in part.

Difficulty conceptually separating the "thing" from various rights, duties, and so on in property relations has also come under scrutiny from another angle. Many theorists have

---

<sup>16</sup> Stephen Buckle (1991, p. 29) describes the *summ* in the natural law tradition as including "what naturally belongs to a person because none of these things can be taken away without injustice." This includes rights to life, bodily integrity, freedoms, and powers and claims over objects of property—the sum or total of entitlements. Accordingly, as Gaus (2021, p. 16) explains, when it comes to distinguishing between a right to free speech and a right to a Subaru Outback, "we should not expect a deep and clear distinction in principle."

observed that the character of “entitlement” is undescribed in the triadic formulation, leaving it wholly parasitic on background notions of an “owner” and a “thing owned.” Earlier I attempted to show the inconsistency in thinking about property as owning an umbrella (as a thing) in some cases and owning some package of use rights in an umbrella in other cases. I pointed out that it is inconsistent to think of the umbrella as the object in the first case and the package of rights as the object in the second. Again, the former is simply a more extensive package of rights than the latter. That led me to suggest that the triadic formulation may be better characterized simply in terms of packages of rights:

*[Henry] holds some rights to [the umbrella] against [everyone else].*

*[Emma] holds some rights to [the umbrella] against [everyone else].*

But what does it mean to ‘hold’ the rights—are they owned? This question requires turning to an ascendent skepticism over the last century which has called into doubt whether ‘ownership’ has any distinctive content at all. Indeed, it is now widely acknowledged that there is no necessary and sufficient cluster of rights that corresponds to the notion of ownership. This acknowledgment was largely driven by legal theorists who observed how the practice of property manifests as disjointed and idiosyncratic in the legal system. Where others saw a unified conception of dominion, these theorists saw an often-unpredictable bundle of rights. For decades, this was one of the most important conclusions of legal thinking about property, though in recent years it has faced some resistance. The following section briefly describes the history of this “disintegration,” as well as its implications for thinking about tax policy, especially as it relates to the legal construction of supposedly pretax property entitlements.

## 1.4.6 The Disintegration of Dominion

The disintegration of ownership emerged out of resistance to the idea that relations among people with respect to resources can be well-characterized by a robust, unified, determinate concept. A fully integrated ownership relation is often described as a kind of absolute dominion; perhaps it is the most extensive package of rights a person can have regarding a resource, consistent with the rights of others. This is sometimes called ‘full-blooded’ ownership, or total dominion, and is typically traced to notions of ownership in Roman law (e.g., Pipes, 1999). The classic example is Blackstone’s description of property as “that sole and despotic dominion which one man claims and exercises over the external things of the world, in total exclusion of the right of any other individual in the universe (1893, Chapter I).”<sup>17</sup> Before Blackstone there was Bartolus de Saxoferrato’s widely cited definition of *dominium* as “the right of complete control over an object, to the extent not prohibited by law” (Rüfner, 2010, p. 127).<sup>18</sup> Bartolus’ definition<sup>19</sup> influenced law for centuries, and was most notably enshrined in the French *Civil Code* of 1804 (*article 544*).<sup>20</sup> Of course even these “purists” recognized limitations on

---

<sup>17</sup> Blackstone was aware that this view lacks any defensible foundation: “And yet there are very few that will give themselves the trouble to consider the original and foundation of this right. Pleased as we are with the possession, we seem afraid to look back to the means by which it was acquired, as if fearful of some defect in our title; or at best we rest satisfied with the decision of the laws in our favour, without examining the reason or authority upon which those laws have been built. We think it enough that our title is derived by the grant of the former proprietor, by descent from our ancestors, or by the last will and testament of the dying owner; not caring to reflect that (accurately and strictly speaking) there is no foundation in nature or in natural law, why a set of words upon parchment should convey the dominion of land: why the son should have a right to exclude his fellow-creatures from a determinate spot of ground, because his father had done so before him: or why the occupier of a particular field or of a jewel, when lying on his death-bed, and no longer able to maintain possession, should be entitled to tell the rest of the world which of them should enjoy it after him. These inquiries, it must be owned, would be useless and even troublesome in common life. **It is well if the mass of mankind will obey the laws when made, without scrutinizing too nicely into the reason for making them**” (1893, Chapter I, emphasis added).

<sup>18</sup> *In primam Digesti Novi partem Commentaria ad D 41.2.17.1 no 4*: “*Quid ergo est dominium? Responde, est ius de re corporali perfecte disponendi, nisi lege prohibeatur.*”

<sup>19</sup> Bartolus also offers a second and broader definition of *dominium*. Property, he says, “may be used to refer in the broadest sense to every incorporeal right, as in ‘I have property in an obligation, for example in a usufruct.’” (*potest appellari largissime pro omni iure incorporali, ut habeo dominium obligationis, utputa usufructus*) (See Boyce, 2007, p. 216)

<sup>20</sup> “Property is the right to enjoy and dispose of things in the most absolute manner provided that the use is not prohibited by laws or regulations.”

absolute dominion such as regulations and so on but, on their view, these amount to merely “refinements outside of the core of property” (MS, 2007, p. 1891).

Though the idea of ownership as total dominion is routinely treated as an omnipresent paradigm, it is in fact something of a historical and cultural anomaly. Contrary to its popular invocation, it cannot be found in the Roman law conception of *dominium*, as Roman property law “has a strong social component” (Boyce, 2007, p. 211). This misunderstanding has its roots in the early modern rebellion of natural law theorists against feudalism, which resulted in a politically useful but ultimately mistaken interpretation of property in Roman law as inviolable (Jones, 1929). Indeed, even the “notion of property as a fundamental right, was a development of medieval and early modern thought” (Boyce, 2007, p. 213). The goal of these early modern thinkers was to end feudalism under which neither feudal lord nor tenant alone could alienate feudal property--dominium utile (or duplex dominium). If divided ownership (*directum* held by the lord and *utile* held by the tenant) was the disease, then perhaps an indivisible conception of ownership was the cure. Dubious interpretations of Roman law were therefore merely means to that end. Nevertheless, the idea of ownership as total dominion went on to influence a number of other legal codes and permeated the common law through Blackstone.

During the twentieth century, however, scholars largely abandoned the idea of ownership as total dominion. Indeed, as the complexity of social relations (including legal relations) saw an enormous increase, many doubted that the idea of ‘ownership’ conveyed any determinate content at all. As an illustration, consider a typical claim about real property, such as ‘Saul owns the parcel 006-577-814.’ What rights and obligations does Saul hold in virtue of ‘owning’ that parcel? One would have to examine a divided structure of governance, including a district municipality (Saanich) which is part of a local government administrative district (the Capital

Regional District) which exists within a provincial region (British Columbia) which is part of a Federation (Canada). Each of these layers of governance contributes to defining the specific rights and obligations that make up Saul's relations to others with respect to 006-577-814. Far from total dominion, Saul's interest will be conditioned by municipal zoning rules which define the permitted lot coverage, number of structures, possibility of subdivision, ratio of basement to non-basement area, and so on.<sup>21</sup> Other municipal bylaws prohibit excessively dirty exterior finishing, unmowed grass in excess of 20cm, and piles of dirt; they protect sightlines, set parking requirements, and limit density.<sup>22</sup> Then the administrative district has its own set of rules, including prohibitions on (among other things) writing on exterior walls and fences, and "offensive or unwholesome matters"; it also authorizes enforcement officers to enter the land "at all reasonable times" to ascertain compliance.<sup>23</sup>

Supposing Saul wants to build a dwelling on the land, that structure must comply with building codes developed by both the federal and provincial governments. These codes require the structure to satisfy thousands of specific rules in demolition, design, and new construction. The rules aim to promote, for example, fire safety, structural safety, safety in use, resistance to unwanted entry, health and sanitation, noise protection, accessibility, and the environment.<sup>24</sup> Provincial rules also define the spatial limits of any interest in the parcel. Medieval scholars understood land ownership as a kind of column of entitlement that stretched from the depths below to the heavens above.<sup>25</sup> *Cuius est solum, eius est usque ad coelum et ad inferos*: whoever's is the soil, it is theirs all the way to Heaven and all the way to Hell. Today, the legal contents of

---

<sup>21</sup> The Corporation of the District of Saanich, Zoning Bylaw 8200, Schedule 220 RS-8 (September 2003)

<sup>22</sup> The Corporation of the District of Saanich, Bylaw 9600, Part 2 (May 2021)

<sup>23</sup> Capital Regional District By-Law NO. 1887 (May 1992)

<sup>24</sup> BC Building Code, Division A, Section 2.2

<sup>25</sup> Michael Heller and James Salzman, *Mine! How the Hidden Rules of Ownership Control Our Lives*, 122

‘full title’ to 006-577-814 do not include entitlements below the surface: “In British Columbia, as in the other provinces, private surface does not include mineral rights. The Government of British Columbia owns and may dispose of subsurface rights to most lands in the province.”<sup>26</sup> And nor does the parcel include entitlement to the heavens, a general rule quickly discovered by those objecting to, for example, construction cranes or the noise pollution of aircraft.<sup>27</sup>

Ownership is also constructed by contract, and therefore laws governing contracts, such as the *Residential Tenancy Act*. Thus we cannot even infer that Saul has the right to enter the owned premises, since the rights to possess and use can be held by someone else, perhaps as part of a fixed term tenancy.<sup>28</sup> Indeed, Saul may both purchase and sell the lot with a tenancy in place, never being entitled to possess the premises at any time in his ‘ownership’ tenure. That possession right may also be divided, so that specific individuals have rights to use specific aspects of the lot at specific times. Or there may be an easement which ‘runs with the land’ such that the ‘owner’ of a neighbouring property has certain use rights to 006-577-814, regardless of who owns either lot. Perhaps these are limited contracts, and ownership corresponds to the holder of the residual interest in the parcel? Even that cannot be right, as the holder of the residual interest is not Saul but rather the state. The state claims an underlying interest which is thought to justify granting land tenure to holders of the land, such as Saul.

If Saul’s title is derivative of the Crown’s underlying title, then it also depends on the legitimacy of the Crown’s assertion of sovereignty. This may seem sound, since the Supreme Court of Canada ruled that “[a]t the time of assertion of European sovereignty, the Crown

---

<sup>26</sup> Government of British Columbia, No. 7, A Guide to Surface and Subsurface Rights and Responsibilities in British Columbia

<sup>27</sup> *Janda Group Holdings Inc. v. Concost Management Inc.*

<sup>28</sup> *Residential Tenancy Act*, 1.1.1

acquired radical or underlying title to all the land in the province.”<sup>29</sup> But even the court admits that title is not absolute; instead it is “burdened by the pre-existing legal rights of Aboriginal people who occupied and used the land prior to European arrival.”<sup>30</sup> Following *Delgamuukw*,<sup>31</sup> the Court interpreted this “Aboriginal interest in land”<sup>32</sup> as the full beneficial interest, implying that “the Crown does not retain a beneficial interest in Aboriginal title land.”<sup>33</sup> *Delgamuukw* also held that Aboriginal title can “compete on an equal footing with other proprietary interests.”<sup>34</sup> In fact the right to “enjoyment of property” is only protected by the *Canadian Bill of Rights* which is a federal statute without constitutional authority.<sup>35</sup> In contrast, Aboriginal title is protected by section 35 of the *Constitution Act, 1982*.<sup>36</sup> Moreover, Aboriginal title to 006-577-814 may exist, and Crown recognition of that title may redefine Saul’s property interest. The Court burdens Indigenous peoples with proving their title, in accordance with three criteria: sufficiency, continuity, and exclusivity.<sup>37</sup> Yet these may be much easier to satisfy than is commonly recognized. If an Indigenous community can demonstrate their exclusive occupancy of the land prior to the assertion of European sovereignty, then the criteria are satisfied. Continuity between present and past occupation is only required when present occupation is relied on as evidence for occupation prior to the assertion of sovereignty. Should Aboriginal title be established, then the Crown’s grant of beneficial interest to Saul infringes on that title. That infringement may be legally justified. If not, the details of reconciliation of the private interest with Aboriginal title

---

<sup>29</sup> *Tsilhqot’in Nation v. British Columbia*, 2014 SCC 44

<sup>30</sup> *Tsilhqot’in Nation v. British Columbia*, 2014 SCC 44

<sup>31</sup> *Delgamuukw v. British Columbia*, [1997] 3 S.C.R. 1010

<sup>32</sup> *Tsilhqot’in Nation v. British Columbia*, 2014 SCC 44

<sup>33</sup> *Tsilhqot’in Nation v. British Columbia*, 2014 SCC 44

<sup>34</sup> *Delgamuukw v. British Columbia*, [1997] 3 S.C.R. 1010

<sup>35</sup> *Canadian Bill of Rights*, S.C. 1960, c. 44

<sup>36</sup> *Constitution Act, 1982*

<sup>37</sup> *Tsilhqot’in Nation v. British Columbia*, 2014 SCC 44

are uncertain. Remedy may include compensation, or it could mean turning the definition and protection of the private interest over to the relevant Indigenous legal authority (Borrows, 2015).

In summary, the applicable rules, and so the contents of ‘ownership,’ can change quite dramatically, even between adjacent parcels, as municipalities may have hundreds of different zones, each with their own specific rules. Thus, the claim ‘Saul owns 006-577-814’ could have a variety of meanings, depending on the governing rules. It should be kept in mind that real property is considered a paradigm case—a simple and intuitive case—of ownership. Other cases are far more complex. The upshot of all this is that ownership claims can be remarkably indeterminate as to who (individual, corporation, partnership, crown) has which entitlements (possess, use, income, exchange) with regard to which thing (land, a pencil, an idea, debt) against who (other individuals, the world, the government). This brief gesture at the litany of rules defining interests in real property suggests that if they are only “refinements outside of the core of property” (MS, 2007, p. 1891) then they are significant refinements indeed.<sup>38</sup>

#### **1.4.7 Disintegration vs. Unified Core**

By the early 1900s, the above-described realities of legal practice began to sway theorists against the possibility of unified, determinate conceptions of property and ownership. The ostensibly traditional view was that “property was understood as entailing a thing owned, and ownership as entailing the owner’s dominion over the thing” (Klein & Robison, 2011, p. 194; *but see* Ripstein, 2011). Such ‘traditional’ theories seemed inadequate as each of the core concepts

---

<sup>38</sup> Other important provincial governance (e.g., in British Columbia) includes environmental legislation such as the *Riparian Areas Protection Act*, the *Expropriation Act*, as well as rules of adverse possession (See the *Limitation Act*, Section 28, in contrast with Alberta which has no prohibiting statute). The federal government also claims powers of eminent domain in the federal *Expropriation Act*. In addition, the meaning of Saul’s ownership of 006-577-814 may also be changed and obscured by other legal arrangements, such as the trust.

(“the thing,” “the owner,” “dominion”) is, at minimum, much more complicated in practice than appearances suggest. As for dominion, it is nowhere found in the institutions of property, and property litigation rarely even refers to ownership, instead centering on the relative strength of claims about the parties’ particular rights and duties (Waldron, 1990, p. 48). Rather than any kind of robust unity, attention to the law reveals “a dizzying diversity of property entitlements” (Breakey). Even the language of “entitlements” and “rights” can be misleading here, as property is “normatively heterogenous” including “not just rights but liberties, powers, liabilities, and immunities as well” (2002, p. 771). As for the owner, the notion of a singular “owner” is similarly misleading given the modern reality of divided control over resources subject to extensive regulation. A theory of property must allow that “there are things that are property but that may not be *owned by* or do not *belong to* a single person” (Wenar, 1997, p. 1944). As for the thing, the “things” referred to in property relations are similarly divided and often abstract. Rights to future income from a legally defined entity like a corporation, for example, can be valuable property but are difficult to even analogize to material objects. All of this heterogeneity and indeterminacy of practice prompted many to doubt that a conceptual analysis of property and ownership was even possible.

Before long, “most jurists” acknowledged the failure of the analytic project, agreeing that “[w]e cannot say that a person owns a resource if and only if he has certain specified rights, powers, liberties, and duties” (Waldron, 1990, p. 47). In its place arose what is sometimes called the bundle theory of property. Bundle theory has been variously described as a disintegration, a deflation, or an elimination of ownership. It accepts that ‘ownership’ is ambiguous and indeterminate and that no particular rights are necessary or sufficient to describe a specific

relation as ownership (Grey, 1980). Instead of determining the ‘owner’ and then inferring rights and duties, the process runs in the opposite direction.

[T]he term ‘ownership’ serves only as an indication that some legal relations, some rights, liberties, powers, etc., are in question. On [the eliminativist] view, the term does not convey any determinate idea of what these legal relations are. In every case, we have to push the words ‘ownership’ and ‘private property’ aside and look to the detail of the real legal relations involved in the given situation (Waldron, 1990, p. 29).

Joseph Singer (1988) provides a well-known application of this kind of analysis to *Local 1330, United Steel Workers v. U.S. Steel Corp.* He explains (p. 621) that the appellate court “wrongly defined the issue as a search for the ‘owner’ of the property. They then assumed that, in the absence of specific doctrinal exceptions to the contrary, owners are allowed to do whatever they want with their property.” Yet searching for the ‘owner’ “is fundamentally wrong. It is simply not the right question. To assume that we can know who property owners are, and to assume that once we have identified them their rights follow as a matter of course, is to assume what needs to be decided” (p. 637-638). The court must “decide who wins the dispute on grounds of policy and morality, and then . . . call that person the owner” (p. 638).<sup>39</sup> On this view, concepts of property and ownership are something like after-the-fact political honorifics bestowed on bundles of rights with no necessary or sufficient content, lacking any unifying idea or normative core. To understand ‘ownership’ in any case we must turn to “the detailed rights that particular people have to do certain things with certain objects, rights which vary considerably from case to case, from object to object, and from legal system to legal system” (Waldron, 1990, p. 26).

It would be hard to overstate the influence of this view, and it became legal orthodoxy for nearly all of the twentieth century. The earliest known usage of the “bundle” term comes from John Lewis who remarked in 1888 that “the dullest individual among the people knows and

---

<sup>39</sup> See also (Glackin, 2014, p. 6)

understands that his *property* in anything is a bundle of rights” (p. 43). By 1945 even the United States Supreme Court disparaged the “vulgar and untechnical sense of the physical thing with respect to which the citizen exercises rights recognized by law” compared with “a more accurate sense to denote the group of rights inhering in the citizen’s relation to the physical thing” (*United States v. General Motors*). The demise of ownership was also the fundamental lesson of property law as it was taught in schools. In 1977 constitutional scholar Bruce Ackerman remarked that “if there is anything a lawyer remembers from his legal education, it is that laymen are deeply confused in their property talk; that the law of property concerns itself with bundles of user-rights, not with the awkward idea that things ‘belong to’ particular people” (p. 115). More than a century after Lewis, David Lametti’s review found that most analyses of property conclude with accepting the bundle theory (1998, p. 666), and Barbara Fried described the integrated view of ownership as “mystifying, a relic of the conceptualism that was the chief casualty of the Legal Realist revolution in the first half of the twentieth century” (2004, p. 72). The dominance of the bundle theory was nearly complete.

Though bundle theorists deny that ‘ownership’ has any definite content, they still aim to provide an analytic framework for thinking about property. One common framework draws from the work of Anthony Honoré and Wesley Hohfeld and so has been named the “Hohfeld-Honoré Synthesis.” Honoré’s primary contribution was an articulation of eleven standard “incidents” of ownership (1961). These incidents characterize the relevant domain—the kinds of relations that are generally called property, such as use, exclusion, income, and so on. None of these incidents, whether alone or jointly, is considered sufficient or necessary to capture the broad spectrum of human relations to which the term “ownership” is applied. But they are standard in the sense that property relations will normally reflect at least some of them.

Later theorists applied Wesley Hohfeld's atomic analysis of rights to these incidents (1917). Hohfeld's work (contra Honoré) also emphasized that property rights are, in the end, about people rather than things. He insisted, that is, that "all rights *in rem* are against persons" (1966, p. 75). He cited Justice Holmes who observed that "the *res* . . . need not even be a tangible thing at all" (Tyler, 1900). Thus, although there are no necessary incidents, property rights do have a general character. They are *in rem* (or multital) rather than *in personam* (or paucital). This distinction, also from Roman law, is one "without which we cannot identify something as pertaining to a distinct legal relation that we may reasonably call 'property'" (Benson, 2002, p. 758). A right *in rem* is often described as "against the world" or against "a large and indefinite class of people" (Hohfeld, 1919, p. 718). In contrast, a right *in personam* is "one of a few fundamentally similar, yet separate, rights availing respectively against a few definite persons" (Hohfeld, 1919, p. 718). A classic example is a promise (or contract) which creates an *in personam* right in the promisee against the promisor. In contrast, a right to exclusive possession of the family home is *in rem*, empowering the holder to expel all intruders (subject to the relevant exceptions). Thus, the bundle theory offers some structure for thinking about property rights, though it rejects the analytic project of identifying necessary and sufficient conditions.

The bundle theory of property forever neutralized the common view that understanding property in practice involves identifying an owner and then conceptually deriving some set of property rights. But it has its detractors, not least because it maintains that the concepts of property and ownership imply little more than a "normative situation" (Ripstein, 2009, p. 24). This left some theorists dissatisfied, and in recent decades this dissatisfaction has grown into a major theoretical project aiming to show that these concepts are more integrated and more useful than bundle theory allows (see, e.g. Merrill 1998, 2001, 2012; Attas, 2005; Claeys, 2011). For

example, Merrill and Smith think the bundle theory fails to account for the role of property as a tool for social coordination. They explain that “The [bundle] metaphor implies that the content of property rights continually mutates from one context to the next as legislatures and courts add new sticks to the bundle and take others out” (MS, 2007, p. 1851). Yet for property to coordinate action, it must communicate rights to a large and indefinite group of potential violators (2007, p. 1850).

“[i]f everyone in the world is expected to respect an owner's right to Blackacre, the content of that right cannot be too complicated or idiosyncratic without placing a large burden on many third parties” (2003, p. 1108).

Contrary to bundle theory, this requires a limited number of simple and easily comprehended rules.

In fact, even Honore allowed that there is a remarkable consistency among different jurisdictions, at least with respect to chattels. As a matter of description, Canadians traveling to Italy need not consult a complex list of juridical rules in order to determine their relations to an umbrella purchased at the Italian corner store. Prescriptively, if knowing someone owns an umbrella fails to convey any content about how we should all relate to that umbrella, then we cannot well coordinate our actions. Some have argued (though this remains controversial (see, e.g., Munzer, 2011; Glackin, 2014; Baron, 2014)) that “property theory in political and legal philosophy has mostly turned away from this deflationary ‘bundle theory’” (Breakey, 2014, p. 585). These views turn instead to what they see as a core “of use, exclusion, and disposition” (Breakey, 2014, p. 585). So perhaps there is some core of primary incidents after all, though more complex cases (which includes the majority of the world’s wealth) notably lack the consistency found in chattels.

For my purpose I see no need to take a side between disintegration or unified core theories of property. What I want to emphasize is that concepts like pretax income are legal constructs. That is not to say that no moral property entitlements could exist outside the law or outside taxation. It is only to clarify the deep sense in which the resource descriptions people claim taxes reduce (including pretax income) are unavoidably legal and posttax. They are a result not only of tax law, but of a market that exists as a manifestation of near limitless legal descriptions. That is, the goods and services that underly pretax calculations are themselves legally described. There is no natural market protected by law—law creates and structures the entitlements that are then subject to market exchange (in accordance with further legal rules). The existence of those descriptions depends on taxation, and the justice of those descriptions depends on the justice of the system as a whole, including taxation. Of course, much more must be said in defense of this last point. Here I only mean to indicate how inextricable the practice of property is from complex and contingent legal rules: “in the modern world, the institution of property is everywhere embodied in law” (Harris, 1996/2002, p. 3).

#### **1.4.8 Limitations of the Legal Paradigm**

One might suspect at this point that I have confused the descriptive for the prescriptive. In focusing on the institutional practice of law, I have missed the real concern, which is that the law itself threatens to undermine moral property rights that stand outside and (should) constrain it. Justice Holmes distinguished the philosopher from the jurist along these lines: “[The relation of legal duties] to moral rights [does] not concern us here. [That is] for the philosopher, who approaches the law from without as part of a larger series of human manifestations. The business of the jurist is to make known the content of the law” (1881/2011, p. 195). In focusing on the

content of the law, perhaps I am simply missing the philosopher's point that moral property rules exert normative authority quite independent of law.

This is indeed a view held among some philosophers who are attracted to something like the Blackstonian view, notwithstanding its inconsistency with the realities of property as an institution. On such views, what it *means* to have a property right includes immunity to taxation. For example,

Maximal private ownership is typically taken to include the right to fully manage (to use, and to allow or prohibit others from using); the right to the full income; the right to transfer fully any of these rights through market exchange, inter vivos gift, or bequest; and the right to recover damages if someone violates any of these rights. Redistributive taxation (e.g., of income or wealth) is incompatible with maximal private ownership (Vallentyne, 1998, p. 611)

“When a conception of full ownership is accepted then clearly involuntary taxation is inconsistent with respect for property rights (the Inconsistency Premise)” (Attas, 2005 p. 119)

To be clear, I do not want to deny that property rights of a certain kind can exist without law. It seems clear that law can be wrong about the package of entitlements it assigns with respect to some resource. For example, the legal right to exclude visitors from one's home includes the right to exclude on the basis of characteristics like sex or race. However, I am skeptical that racism and sexism are part of the moral right to exclude. The law can also assign rights to the wrong people, as it does when someone is morally entitled to certain uses of a resource, but someone else is legally entitled to exclude from those uses.

The crucial takeaway for my argument is only that, whatever one thinks about morality and property rights, one cannot stand in the relation of “full ownership” or “maximal private ownership” to some pretax legal construct *in virtue of* it being a legally generated entitlement. Pretax *income* is a politically relevant and useful example, but it is only one of many ostensibly pretax legal constructs to which the same argument against the philosopher's view of maximal

ownership can be applied. The common line of thought looks like this: “The law calculates that I earned \$500,000 but that taxes will take \$200,000. Therefore, that’s \$200,000 of my money that the law is taking from me.” If you “maximally own” that \$500,000, the philosopher might argue, then taxing it violates your property rights. Maybe so, but the point here is that people do not, and cannot, “maximally own” a pretax legal construct. This point is elaborated in detail in chapter three, but here I only want to connect it to the preceding comments on property.

Note that claiming ownership over \$500,000 of pretax income is a moral claim and so requires justification. The moral hurdle for taxation implied by the thought above—“that’s \$200,000 of my money that the law is taking from me”—only exists if we accept entitlement to \$500,000 as a kind of fixed moral point. That fixed point provides a baseline distribution with some assumed moral salience or force, against which the interference of taxation is then measured and evaluated. The distributive baseline therefore morally constrains the realization of tax justice. But why is \$500,000 accepted as a morally significant baseline? The tax objector does not offer a theory of distributive justice under which their<sup>40</sup> entitlement happens to *coincide* with the \$500,000 pretax figure. Instead, they simply point to the fact that \$500,000 was an output of the market-legal system, outputs of which are at least *prima facie* justified. That number was calculated in accordance with the relevant legal instruments that together define its underlying economic character in accordance with a legal definition of income. The “economic character” I have in mind here is just an application of the above-described complexities of property and the other legal aspects of a market economy. For example, if the \$500,000 was calculated in accordance with operating a factory, then it would emerge from a complex interplay of legal standards related to employment, real property, corporations (including limited

---

<sup>40</sup> Throughout this project I intentionally use the singular they.

liability), consumer protections, importing and exporting, and so on. These and other rules construct a market-legal system from which that figure emerges.

This market-legal system can deliver morally justified outcomes (entitlement). But since taxation is an essential part of the system, tax justice is antecedent to the general justice of the system's outcomes. Preliminary outcomes like pretax income cannot stand as fixed moral points *constraining* justice in taxation, since their moral salience is grounded in the justice of the system which constructs them, and that system includes taxation as an essential element. In other words, pretax income could only constrain justice in taxation if tax justice was already achieved. As an outcome of a market-legal system, any moral weight of pretax income must *depend on* justice in taxation and so cannot function as a point of fixed moral significance in realizing that very justice. If the system generating those descriptions lacks justification, then those descriptions do not have moral force in virtue of being outputs of that system. As pointed out above, whether the person with a pretax income of \$500,000 is entitled to that amount, to \$200,000, or to something else altogether, depends on some justificatory story. It is incoherent to claim entitlement to \$500,000 *against taxation* in virtue of the fact that \$500,000 is the output of a market-legal system the outputs of which are justified only if taxation is already just. This is considered more fully later, but the short form is that it is a mistake to appeal to ostensibly pretax legal descriptions as morally fixed points in determining tax justice.

I do not see this as confusing the legal (descriptive) for the philosophical (prescriptive). Instead, it is an argument about the moral status of a legal description. When we talk about owning this or that, we naturally conflate a moral sense of entitlement with a legal description of

power. But there is no reason to take one point<sup>41</sup> (such as pretax income) in the process of calculating legal entitlement as representing some hallowed ground. The way we “own” property in a legal system looks nothing like the philosophical starting point of total dominion. Trying to derive moral implications from that conception and then map them onto legal constructs like pretax income is hopelessly misguided. Indeed, Attas (2006, p.123) argues that something like what I describe here is not an account of property but rather an argument for the impossibility of property.<sup>42</sup>

In one sense he is exactly right. The idea that people in general are “full owners” of pretax legal descriptions of resources is indeed impossible. But that, it seems to me, is simply a reason to doubt the utility of “full ownership” in the tax context. We need a different conception of property, and one that better coheres with its institutional manifestations. Theorists concerned with developing a conception of “full ownership” are not necessarily misguided, but it is a mistake to try and apply that work to concepts like pretax income. If people are fully entitled to resources equivalent to some pretax legal description *regardless* of what the law says (not *because* the law says so) then a separate justificatory story is required. Surely it would be odd to claim that, as it turns out, people are entitled (as a matter of coincidence?) to exactly what the complex array of legal rules determines their pretax income to be. That would be an especially odd claim, since generalizing such an entitlement seems inconsistent with the existence of the

---

<sup>41</sup> “If, as we have postulated, individual rights are defined as rights to do things with respect to some initial set of endowments or goods, along with membership in a collectivity that is empowered to act by less-than-unanimity rules, and, further, if these rights should be mutually accepted, it becomes inconsistent and self-contradictory for a person to claim that his “rights” are violated in the mere working out of the collective decision rules that are constitutionally authorized.” Buchanan, *the limits of liberty*, p. 57

<sup>42</sup> In considering Murphy and Nagel’s argument (2006, p. 123) he writes that they “have appropriated the concept of property to their egalitarian agenda. I have much sympathy for their ideal, but I think they are wrong to represent it in terms of property rather than as the abolition of property.” He goes on to claim they are committed to “a revisionist and certainly unusual account of property.” One of my objectives in Chapter three is to disentangle their rejection of entitlement to a pretax baseline from their conventionalist account of property.

conditions that generate it. That is, full entitlement to pretax income is inconsistent with the existence of the market-legal conditions that generate the resources at issue. In any case, such a theory would need to work *towards* the justification of entitlement the same economic power identified by pretax income rather than taking that entitlement as a given baseline to which some concept of “maximal ownership” can be straightforwardly applied.

#### **1.4.9 The Objects of Ownership: from Self to Wealth**

We have seen that property is a bundle of incidents, even if it has some unified core. In the absence of necessary and sufficient conditions, theorists have turned to paradigm cases and family resemblance to help illuminate property in its multiple and complex presentations. This involves resolving some features of property in simple cases and then testing their application in more complex cases. Today, material objects are taken as the paradigm case with the ambition of working towards other cases like electromagnetic spectra, data, trademarks, fiat currency, patents, use rights (to fish, to mine, etc.), financial instruments (e.g. options, derivatives), corporate shares, reputations, artistic expression, labour, debt, and most controversially, the self. In short, the theoretical direction runs from simple cases of material wealth to the most complex case of the self.

[we must begin with] “rules governing access to and control of material resources” (Waldron, 1990, p. 31).

“If we cannot provide a suitable and coherent account of property in corporeal objects, we cannot hope to do so for other more difficult instances, such as intellectual property” (Benson, 2002, fn. 15).

The core idea of property is the ownership of some material object, such as a car or a violin. But the idea is often extended to include money, patents, undeveloped resources, shares, future payments, abilities, options, ideas, reputation, and so forth. I will restrict the discussion of property to a relatively narrow sense, to include only material objects and money. The issue of justification is difficult enough, and I want

to avoid complications that follow from a consideration of intangible possessions. If the justification works, it may be extended beyond the narrow sense (Kekes, 2010, p. 3)

“property rights relate to tangible property in the first instance, land and chattels...” (Penner, 2020, p. 5).

“the family home is perhaps the most obvious paradigm of a property right” (Breakey, 2014, p. 592).

“[T]he ownership of material things . . . is the proper model for the concept of self-ownership” (Cohen, 1995, p. 215).

“Agents are full self-owners just in case they own themselves in the same way they can fully own inanimate objects” (Vallentyne, 2005, p. 202).

This theoretical progression does track a common intuition that we have a better grip on what it means to own something like an umbrella compared with owning a patent. As noted above, there is some uniformity found in cross-cultural expressions of property and ownership when applied to chattels. But it is less clear that starting with this intuition is the most useful way to conceive of property. After all, it seems obvious why this has reinforced the “total dominion” view that, I argue, is deeply misleading in many important cases.

What is the alternative to the material objects paradigm? Admittedly, it makes little sense to start with the hardest and least intuitive cases and work towards the simpler and clearer ones. Nevertheless, I think it is worth considering whether our interpersonally constrained sense of *mine*, of agenda setting, and of authority, really are most clearly and helpfully expressed in cases involving material objects. Though the modern consensus clearly sees material objects as the core or paradigm cases of property, this was not always the case. As noted above, there is an older natural law tradition of *summ*: “what naturally belongs to a person because none of these things can be taken away without injustice” (Buckle, 1991, p. 29). James Madison wrote that property in its “larger and juster meaning, [...] embraces every thing to which a man may attach

a value and have a right; and *which leaves to everyone else the like advantage*” (1792). On this view, property is “an extension of the sort of ownership that we have of our own selves” (Ripstein, 2004, p. 13). This represents a kind of inversion of the modern view, in which we work outwards from persons as a paradigm case: from self to wealth.

Arguably, Locke followed this tradition when he described individuals uniting in society “for the mutual preservation of their lives, liberties and estates, which I call by the general name, property” (1690, IX.123). Property was not related to material objects in the first instance, but rather the self, and only extended from there. Similarly, Mill seems to accept incorporeal faculties as a paradigm case, and from there moves to corporeal objects:

Nothing is implied in property but the right of each to his (or her) own faculties, to what he can produce by them, and to whatever he can get for them in a fair market; together with his right to give this to any other person if he chooses, and the right of that other to receive and enjoy it (1870, II.II.3).

Even the previously mentioned distinction of *meum* and *tuum*, fundamental to social organization, surely originates in the body and faculties as much as in the external world. Our early characterizations of property—agenda setting, final determination of use, special authority, freedom within a constrained set of options—all seem apt descriptions of our relation to ourselves.

Thinking of selves in this way raises some well-trodden objections, partly with good reason. The first and often most concerning is that self-ownership has morally deranged logical implications. Arguably, libertarian attempts to work out the logic of self-ownership have offered a kind of moral *reductio ad absurdum*. I have in mind Robert Nozick’s infamous claim that a free system “will allow [an individual] to sell himself into slavery” (ASU, 1974, p. 331). Other moral concerns include that self-ownership in a market environment “runs afoul of an important

and basic moral principle that one ought not to treat others or oneself simply as means to one's ends" (Brenkert, 1998, p. 52). In addition to these moral failings there is also the concern that selves differ from objects of wealth so significantly that the analogy simply fails. This is captured by the common claim that self-ownership is impossible because only things can be owned, and persons are not things (Russell, 2018). Perhaps it is worth noting that G. A. Cohen (1995, p. 212) rejected this argument as question begging, but nevertheless attributed it to Kant. To be fair, Kant did identify a specific difference:

“someone can be his own master but cannot be the owner of himself (cannot dispose of himself as he pleases)—still less can he dispose of others as he pleases—since he is accountable to the humanity in his own person” (1797, p. 56).

Others have generally followed the same path, identifying power over alienation as the fundamental sense in which persons differ from things. Ripstein maintains that “strictly speaking, I do not have property in my own person; there are some inherent limits on my possessory rights, so that, for example, I may not alienate my own person. But I have rights in my person like those I have in external things” (2004, p. 13). Similarly, John Kekes argues that we are contingently connected to objects of ownership, unlike the self which cannot be alienated. Thus, he concludes, “it is highly misleading to claim that having a self is like owning a material object” (2010, p. 2).

I think the objectionable moral implications and the concerns about disanalogy arise from conceptions of property and ownership that are inconsistent with property as an institutional practice. Of course, we do not want conceptions of property and ownership that appear to license exploitation, human alienation, domination, abuse, or enslavement. But it seems to me that these concerns simply fail to take seriously the demise of the analytic project to define property in terms of necessary and sufficient conditions. Consider first the claim that we can alienate

material objects but not ourselves. Both Kekes and Ripstein reject the comparison between the self and material objects because the latter can be alienated while the former cannot. That suggests that the power of total alienation is an essential incident in a workable notion of property. But there is no theoretical requirement to accept alienation as a necessary condition of a property relation. Certainly, in practice such a requirement is routinely violated. As with all incidents of property, there is a long history of regulation and outright prohibition on alienation. Aboriginal title is an obvious and modern Canadian example, but there is also the English history of subinfeudation and the general practice of personal licenses and permits which clearly feature “inherent limits on. . .possessory rights.” Even Aristotle reports the practice of the Locrians who, on the island of Leucas, were forbidden by law to sell personal property without proof of misfortune (Politics, Book II, Part VII). In short, there is no conceptual contradiction in the claim that some material object is my property but that I may not alienate it. What that claim (“it is my property”) *means* must simply be worked out in context. Presumably it implies an *in rem* character of several of the standard incidents of property, with no requirement that total power over alienation be among them. There is “no such thing as a generic property right” (Pistor, 2019, p. 126), and therefore no such thing as a “straightforward application of the concept of ownership to the case of x” (Narveson, 1998, p. 8).

What I want to suggest is that the limitations and context-dependence we see in our powers over ourselves are not reasons to reject the self as a paradigm but precisely the opposite. Powers to exclude, to use, and to set the agenda are found most profoundly in the body and self rather than the external world. The fact that these powers are limited by moral considerations illuminates the indeterminacy of property in material objects, rather than the fiction of total dominion over the world telling us something morally repugnant about how we relate to

ourselves and each other. Once we have set that fiction aside, it is obvious that concerns about people selling themselves or others into slavery simply do not follow from a workable conception of property. Consider that the wrongness of slavery is that a slave is “always subject to the will, or grace, of his or her master” (Ripstein, 2004, p. 8). Therefore, our conception of property simply excludes this possibility. As we have seen, real control over resources (like selves) is always severely limited in comparison to an abstract notion of “maximal control.” Of course, people are not *only* resources, but they *are* resources. When too much control leads to unwelcome outcomes, the extent of control is modified. Widespread constraints on dead hand control and the rule against perpetuities are well-known, but they are only examples of myriad “limitations” designed to achieve various goals. Locke himself argued that property excludes the right to destroy: “Nothing was made by God for Man to spoil or destroy” (1960, s. 31). To his point, suppose a wealthy person buys up valuable art or other cultural artifacts simply to set them on fire. Is permission to do so built into the very idea of property? I suggest there is no reason to assume it must be. In short, there are no conceptual implications of something being “property”; meaning is a matter of policy.

Once we have set aside the notion of total dominion and the notion that any particular right, duty, liability, or power is inherently lodged in our conception of property, then we are free to simply shape the conception as it fits the context. As Shane Glackin puts it, “[n]o moral or legal conclusion can or should be drawn from the declaration that some thing is property or that some person is its owner” (2014, p. 2). This, of course, is already what happens as property relations are expressed through institutions. Mill’s consideration of inheritance (though not of bequest) captures this way of thinking about property.

That the property of persons who have made no disposition of it during their lifetime, should pass first to their children, and failing them, to the nearest relations, may be a

proper arrangement or not, but is no consequence of the principle of private property. Although there belong to the decision of such questions many considerations besides those of political economy, **it is not foreign to the plan of this work to suggest, for the judgment of thinkers, the view of them which most recommends itself to the writer's mind** (Mill, 1909, II.II.3, emphasis added).

Entitlement to oneself is a way of capturing what it means to be free (to set and pursue one's own ends). As Narveson puts it, "To say that one "owns oneself" is to say that others may not use one's *self* without that self's permission. Why, then, would there be a problem?" (2010, p. 103). This characterization of the way individuals relate to themselves as resources strikes me as accurate and inoffensive. The problem comes downstream with the assumption that ownership *also* implies other things like a right to total alienation, which implies the permissibility of slavery, and so on. But on my view that and similar logical chains are simply non sequiturs. The incidents of property and so the nature of ownership must be justified in context, and different contexts will permit different justifications.

"the right to private property is a bundle of rights and the components of the bundle need to be justified differently. These components are the *possession, use, and limits* of private property. The second reason is that the various rights to private property must be justified differently in different contexts. It is highly unlikely that property's justified possession, use, and limits would be the same in feudal, revolutionary, post-Communist, ex-colonial, tribal, impoverished, affluent, agrarian, industrial, authoritarian, and democratic societies" (Kekes, 2010, p. 14)

In the case of human powers as resources, some exclusion and use may be justified, while total alienation is not. There is no contradiction with a background conception of total dominion because such a conception is rejected from the outset.

This more open-ended way of thinking about property and ownership reflects what is seen in legal practice and is therefore directly relevant to how we conceive of the legal construction of pretax holdings. Legal descriptions of pretax holdings do not track some logical derivation from initial holdings, through market exchanges, finally to arrive at a morally justified

distribution awaiting the interference of taxation. Instead, the legal system applies a contingent conception of income to an unimaginably large variety of rules that characterize market transactions in a specific way. But that characterization before some tax or another is included in the calculation reflects a much looser notion of ownership than the commonsense of total dominion. It is contingently designed to achieve certain ends, and so it does not arrive with a built-in moral justification. Some argument is necessary to explain why people would be entitled to that particular step in the legal process of defining legal entitlements. Arguments that *begin* with total dominion over that step simply miss the point. In my view, the institutional expression of property has much more to teach us about how we should conceive of ownership than rigid philosophical accounts of ownership have to teach us about property in practice. At minimum, it is unclear why some rigid moral concept of dominion would map directly onto the fluid and contingent legal construct of pretax income.

### **1.5 Chapter 1 Conclusion**

As stated in the introduction, the overarching goal of this project is to reveal how conceptions of fairness and efficiency in tax justice can be (and often are) corrupted by affording moral significance to arbitrary baselines. Conceptions of property and ownership are inextricably linked to this project because ownership of certain property prior to taxation is the primary baseline against which taxation is criticized as interference. This chapter provides some of the necessary background for my later arguments against the moral significance of that baseline.

I began this chapter by framing property as a social strategy for coordinating resource use. This suggests that property rules are malleable conventions constructed in service of other values. In contrast, I introduced the view that property rights are morally fundamental and exist

anterior to civil society. Arguments about tax justice normally afford great significance to this debate. The thought is that accepting a conventionalist view implies that pretax descriptions have no special moral status, while a rights view suggests the opposite. My view is that this framing obscures the crucial fact that a rights view does not automatically justify a general entitlement to pretax descriptions. Further, in chapter three I argue that it *cannot*.

To defend my view, I traced some of the debates about the meanings of property and ownership. My path through these debates was chosen to highlight several key points. The first was that “property is not simply a series of propositions about chattels or land but describes a realm of social relations” (Bryan, 2000, p. 7). Those relations are commonly understood in triadic form as A owns B against C. This is the source of some confusion, since both B and the triadic relation as a whole are called property. From Roman Law until today, attempts have been made to resolve this confusion by distinguishing property *rights* from the *objects* of those rights. Property is *tangible* while rights are *intangible*, or B *is* property while the triadic relation captures property rights *in* B. I argue that such attempts fail for the same reason: property rights and their objects are not mutually exclusive categories. Property always involves standing in relation to others with respect to some resource, but that resource may itself be rights.

Disagreement about the nature of property is compounded by controversy about the meaning of ownership. In the institutional context, the notion of ownership has minimal significance. Instead, conceiving of property as a bundle of rights has dominated legal theory and practice for over a century. There is debate over the malleability of the bundle, but it is widely accepted that the concepts of property and ownership have no necessary and sufficient conditions. There are standard incidents, and they are typically *in rem*, but there is no general practice of identifying owners and conceptually deriving entitlements. Instead, the rights, duties,

liabilities, and so on are constructed in context as a matter of policy. The result is heterogenous, dynamic, and at times idiosyncratic. Certainly, it looks nothing like the philosophical ideals of “maximal” or “full-blooded” ownership.

I considered whether the preceding conclusion simply misses the point. Perhaps if legal practice falls so short of a philosophical ideal, then so much the worse for the moral status of that practice. Maybe so, but that objection misses what I see as the relevant point for tax justice. My argument is that concepts like pretax income *are* legal constructs. So, applying concepts like “maximal ownership” to such constructs as if they are fixed points of moral significance is a mistake. It leaves unanswered *why* people would be entitled to the figure calculated by a particular step in “the mere working out of the collective decision rules that are constitutionally authorized” (Buchanan, p. 57).

The final section of this chapter suggests we rethink our commitment to the “material objects” paradigm of ownership. That paradigm suggests that ownership is more uniform and more consistent than it is in the most important cases. It also reinforces the total dominion view, since that view is least problematic in cases of chattels such as umbrellas or toothbrushes. But when we work out from there, we find that complex cases of abstract property relations, including the way we relate to ourselves, look remarkably dissimilar to cases of material objects. I suggest that the way we think of ourselves may be a better model for thinking about other kinds of social relations to resources. The self is where we find most profoundly a sense of authorship or of agenda-setting. Yet there are important moral limits to how we treat ourselves as resources. The same is true of umbrellas and patents; as the significance and context differ, so should the permissible property relations. I think this offers a more useful starting point than the tired claim that ownership implies total dominion and therefore some other set of morally disastrous

implications. I do not think that is an illuminating way to think about ownership, and most importantly I see no reason to think people own their pretax income in that maximal sense. As a result, I see no reason to think about tax justice in terms of state interference with individuals' total dominion over pretax holdings. Moreover, as I argue in chapter three, I doubt such entitlement is even conceptually possible. Having considered the core background concepts of property and ownership, I can now turn to attempts to theorize tax justice directly. Chapter two outlines prominent approaches to the normative dimensions of taxation, as well as my assessment of their strengths and weaknesses.

## **Chapter 2: Normative Taxation**

### **2.1 Chapter 2 Introduction**

Taxation is perhaps the most important mechanism for realizing a conception of distributive justice. It also confronts citizens with the coercive power of the state in an immediate way. Yet there exists no widely accepted theory of tax justice. This is partly explained by the protean character of modern taxation: taxes allocate resources, create incentives, fund public goods, address collective action problems, and more. As well, claims about fair taxation always implicate technical and practical considerations alongside their normative dimensions.

Historically experts in the technical and practical (such as economists and policy makers) have more readily engaged this tangle of considerations than experts in normative theory (such as philosophers), though that is beginning to change. The results of this engagement are fragmentary and often inconsistent. However, in this chapter I argue that the fragments can be roughly sorted into two broad approaches to questions of tax justice.

The first approach assesses taxation as an institutional interference with a pretax allocation of resource entitlements. It conceives of the collective tax burden as a social invoice that must be fairly distributed across that pretax allocation. Thus follows various principles of distribution: the tax burden should be distributed according to ability to pay, or benefits received, and so on. But the second approach argues that this project of fairly distributing the tax burden is misconceived for two reasons. First, it is myopic in its assessment of particular taxes without considering how those taxes fit within the broader institutional arrangement. Second, it presumes an existing allocation of resource entitlements with which taxes interfere. In a modern state, however, taxes are antecedent to and so already presumed by any allocation of entitlements. Instead of attending to a fair distribution of an illusory tax burden, this second approach

conceives of taxation as constructive social architecture. Accordingly, it holds that taxes should be assessed in terms of their contribution to a distribution that satisfies the appropriate principles of justice, whatever those principles may be.

## **2.2 Taxation as a Normative Inquiry**

Few public debates are as polarized as taxation. Taxes are at once celebrated as crucial for large-scale cooperation (Piketty, 2014) and denounced as “on a par with forced labour” (Nozick, 1974, p. 169). Disagreements can be technical, involving sophisticated predictions about the consequences of various proposals for the economy as a whole. But in the end “[t]axation is not a technical matter. It is preeminently a political and philosophical issue, perhaps the most important of all political issues” (Piketty, 2014, p. 630). This irreducible normativity is widely accepted, and most familiar debates involve appeals to fairness from all sides. For instance, when proponents of a wealth tax argue that it will compel the rich to pay their fair share, opponents answer that it will unfairly deprive people of their rightful earnings. Or a carbon tax is proposed to reduce emissions, but critics claim it will unfairly impact people with low incomes. Similarly, there are disputes about whether tax structures that seem to favour certain preferences are unfair. For example, tax deductions might arbitrarily favour preferences to donate to charity, to own rather than rent, or to save rather than consume. Consequently, the overall system of taxation is sometimes criticized as unfair for failing to provide “neutral tax treatment of economic activities” (Slemrod & Bakija, 2017, p. 228). And the list goes on.

The ubiquity of such debates calls for a principled account of tax fairness. Many have remarked with some surprise that philosophers have largely ignored that call (e.g., O’Neill & Orr, 2018, p. 1). Instead, existing accounts of tax justice come largely from economics, law, and

public policy. In recent years, however, taxation has received increased attention from philosophers. The resulting literature is complex, notably fragmented (Halliday, 2013), and often guided by implicit assumptions about the nature of property rights. However, this chapter identifies a certain structure in the fragments, and so roughly cleaves the discussion into two general approaches: tax burden analysis and tax as social architecture.

Tax burden analysis has a long history (Seligman, 1908), but it was elaborated in detail by various influential public finance economists in the 20<sup>th</sup> century. The approach relies on a particular view of the relationship between the market and the state. It conceives of taxation as an interference with a distribution of entitlements generated by private economic activity. That interference is held to impose a burden both collective and individual. On the collective side, taxes “unquestionably impose costs on the economy” through their distortion of patterns of private exchange (Slemrod & Bakija, 2017, p. 227). And the individual costs seem easy to see— simply compare individual holdings before and after tax. For burden analysis, then, taxes are characterized as the coercively imposed surrender of private gains to, as Oliver Wendell Holmes famously put it, “pay for civilized society” (*Compañía v. Collector*, 1927).

Tax burden analysis pairs that portrayal of citizen and state with an intuitive normative commitment: a fair distribution of the costs of civilization must treat equally situated people the same, but also display sensitivity to relevant differences among taxpayers. Two primary accounts have been developed to elucidate the pertinent norms of equality and difference: “the theory of benefits [and] the theory of ability” (Seligman, 1908, p. 148). Both theories articulate these norms with reference to a pretax distribution upon which taxes are imposed. The theory of benefits holds that an individual’s tax burden should be imposed in proportion to the benefits she receives. In contrast, the theory of ability holds that her tax burden should scale with her ability

to pay. These measures are then used by the tax burden analyst to assess differences in holdings resulting from various tax rules and finally to ground judgments of fairness.

A strikingly divergent conception of taxation as social architecture motivates the second approach to tax fairness. This conception is found both in the philosophical arguments provided most prominently by Murphy and Nagel in *The Myth of Ownership* (2002) and in the mainstream economic theory of optimal taxation (Fleurbaey, 2018). The social architecture approach is distinguished by its rejection of the traditional conception of taxation as an interference with a pattern of pretax entitlements. Instead, it characterizes taxation as a constructive and institutionally embedded technology. The basic contours of this approach are best illuminated by considering the two complementary challenges it poses to tax burden analysis.

The first challenge is that tax burden analysis fails to properly regard taxation as a technology embedded in a broad scheme of institutions. It assumes that principles of fairness can guide equitable distribution of the tax burden. But such narrow judgments of tax fairness are criticized as myopic for their disregard of the overall system of tax and *transfer* in which all tax burdens are situated. Moreover, that tax and transfer system is itself embedded, and so cannot be assessed independently from the comprehensive set of social institutions of which taxation is but one essential part.

The second challenge is that attempts to measure the burden of taxation illicitly countenance an unjustified allocation of resources. But the state—and thus taxation—is a precondition of any ostensibly pretax allocation of resources. So the central postulate of tax burden analysis, the pretax baseline, turns out to rest on the “conceptually invalid” illusion that taxation and the outputs of the private economy are somehow dissociable (Shoup, 1969, p. 577). Rejecting this illusion means that a tax burden cannot be understood as an interference with a

pretax distribution of resources. Without a pretax distribution, the analysis loses its baseline against which to measure the tax burden. Put another way, tax burden analysis errs by accepting a set of pretax entitlements as an independent normative standard against which to evaluate the tax system without recognizing the essential role of taxation in constructing that very set of entitlements.

In combination these two challenges invite skepticism about the possibility of developing narrow principles of tax fairness that govern distribution of the tax burden. By refuting the fundamental assumptions of burden analysis the challenges force “a reassessment of traditional tax policy approaches” (Sugin, 2004, p. 1992). In turn, however, they make room for an expansive account of taxation as social architecture. Such an account rejects a construal of institutions of governance as bills to be paid, mere parasites on a private economy. Instead, those institutions are seen as productive actors that create and shape markets, products, and citizens themselves. Denying the normative authority and even coherence (see chapter three) of pretax resource allocations bolsters this instrumental, architectural view of taxation. Accordingly, tax is conceived as a protean technology for constructing an institutional regime consistent with principles of justice. Taxes raise revenue, configure markets, curb the money supply, allocate resources, structure incentives, and so on. Instead of attending to the fair distribution of the tax burden, therefore, normative inquiry must examine the comprehensive fairness of post-tax allocations: “outcomes, not burdens” (Murphy & Nagel, 2002, p. 98). In some measure mainstream academic economics has already made this transition, abandoning tax burden analysis in favour of searching for parameters that optimize a given social welfare function (Fleurbaey, 2018). Still, both public policy and public opinion often cling to the moral salience of pretax allocations and so to the promise of an equitably distributed tax burden against that

baseline. Thus, an overview of tax equity must yet begin with burden analysis and the public finance orthodoxy.

### **2.3 Tax Burden Analysis**

The orthodox public finance approach to tax equity was first elaborated in the works of Arthur Pigou (1947), Henry Simons (1938 & 1946), Robert Haig (1921), and Richard Musgrave (1959 & 1989). This approach accepts an exogenously determined tax burden that must be distributed among taxpayers and aims to develop equitable principles for its distribution. Assessments of tax equity proceed by applying those equity principles to compare pre- and post-tax worlds.

A pretax world provides a baseline against which some proposed tax burden is measured and assessed, and it can be variously described. For example, the burden of taxation might be assessed in terms of its consequences for the pretax distribution of utility (Feldstein, 1976) or income (King, 1983). That assessment might compare the predicted effects of taxation with a hypothetical tax-free world, or a world that maximizes a social welfare function (Kaplow, 1989). In practice, however, tax proposals are inevitably about reform rather than design. As a result, the existing pre-reform distribution provides a ready baseline, rather than appealing to a hypothetical world (Feldstein, 1976). However, this way of proceeding has been properly criticized for arbitrarily privileging the status quo distribution (Kaplow, 1989).

In evaluating the equity of both design and reform proposals, the orthodox public finance approach begins with the largely uncontroversial view that taxation should not arbitrarily discriminate, but that it should be sensitive to some differences among taxpayers. It captures these commitments in tax specific terms as horizontal and vertical equity (Musgrave, 1990).

### 2.3.1 Horizontal and Vertical Equity

The idea that tax justice compels some kind of equal treatment is an old one (Elkins, 2006), but it was expressly applied to the tax burden by Henry Simons who asserted that “taxes should bear similarly upon all people in similar circumstances” (1946, p. 8). A decade later Richard Musgrave coined the term “horizontal equity” and affirmed it as “perhaps the most widely accepted principle of equity in taxation” (1959, p. 160). Musgrave attributed its wide acceptance to its “reflecting the basic principle of equal worth” and offering “protection against arbitrary discrimination” (1990, p. 113). Tax policy analysts ordinarily integrate considerations of horizontal equity with its counterpart, vertical equity, which requires “appropriate differentiation among unequals” (Musgrave, 1990, p. 113).

Application of the two principles can be illustrated with a simple case. Suppose two individuals have equivalent pretax holdings (income, consumption, wealth, or some other tax base). A horizontally equitable income tax (for example) will require the same tax payment from both individuals. Given their equivalent pretax position, different tax liabilities raise concerns of inequitable treatment. For instance, suppose the state allows taxpayers to deduct mortgage interest payments from taxable income. Now individuals with pretax holdings of equivalent value may nevertheless owe different amounts of tax, depending on the arbitrary fact of whether or not their holdings include a mortgage. According to the public finance view, this difference in tax burdens is a violation of horizontal equity. In complement, vertical equity applies when pretax holdings are not equivalent. In such cases higher payments are required from those with greater pretax holdings. Here concerns of inequity arise when individuals with higher pretax holdings pay equal or lower tax compared to those with fewer holdings.

Though policy analysts normally rely on both horizontal and vertical equity, there is some debate in the public finance literature about how these principles are related. Musgrave originally described them as complementary, simply “different sides of the same coin” (1959, p. 160), though he later recanted, identifying an independent role for horizontal equity when evaluating second-best solutions (1990, p. 114). Others construe vertical equity as “analytically more fundamental” (Murphy & Nagel, 2002, p. 13) and horizontal equity as adding nothing substantive, merely a “measure in search of a principle” (Kaplow, 1989). Still others have gone so far as to declare the two principles “fundamentally incompatible,” given their final grounding in conflicting strains of political philosophy (Elkins, 2006, p. 88). Such tensions tend to emerge when trying to specify consistent arguments that secure both principles. For example, vertical equity is often defended as welfare maximizing, and so as part of the social welfare tradition (Kaplow, 1989). In contrast, horizontal equity maintains the status quo regardless of welfare consequences, clearly conflicting with the social welfare tradition.

Whatever their relationship, all agree that these principles provide limited guidance in their “barest form” (Pigou, 1947, p. 5). Musgrave, for instance, insisted that “measures of tax equity must be grounded in a view of entitlement and distributive justice” (Musgrave, 1989, p. 114). This would include some account of fair discrimination by identifying relevant differences for the purpose of taxation. Edwin Seligman called the provision of such an account “the fundamental question as to the basis of taxation” and he argued that there are two contenders: “the theory of the benefits or the theory of ability” (1908, p. 148).

### **2.3.2 Ability to Pay**

The principle that individuals “should contribute to the public burdens in proportion to ability” was developed in the middle ages, and its general outline can be traced to the Greeks (Seligman, 1908, p. 205). It remains deeply influential in modern accounts of tax fairness (Slemrod & Bakija, 2017). Controversy persists over the proper index of the ability to pay. Ability might be gauged in terms of resources, income, welfare, or some other standard. In addition, given some index, the tax burden could be calibrated to either its potential or actual level. Plainly, an individual’s potential income could greatly exceed her actual income. Efforts to capture that potential in the tax base motivate an endowment tax.

#### **2.3.2.1 Endowment**

Endowment is capacity to generate or acquire resources (or welfare, etc.). Historically this mainly referred to earnings generated by real property and other forms of material wealth (Seligman, 1908, p. 206). But its meaning transformed with the ascent of human capital in the hierarchy of productive capacity. It is now well-recognized that potential to earn is modulated by knowledge and expertise, personality, connections, credentials, and so on. (Murphy & Nagel, 2002, p. 21). Moreover, human capital is thought to be “most individuals’ primary asset” (Kaplou, 1994, p. 1478).

Even in principle, taxing capacity is a double-edged sword. On one side are the economic benefits. People subject to an endowment tax are given an incentive to act productively, such that their actual earnings best realize their capacity. On the other side is concern about erosion of autonomy. Sufficiently high levels of taxation could blur the line between fair incentives and unjust coercion, raising the Nozickian specter of “forced labour” (Nozick, 1974, p. 169). The

closer taxation approximates the economic ideal of maximum productivity, the greater the threat to autonomy.

The economic argument in favour of taxing endowment is that it would create an income effect and no substitution effect. An endowment tax has no substitution effect because it does not change the relative costs of labour and leisure: it assigns the same tax liability regardless of how much an individual works. In contrast, a tax on labour income decreases the reward for work and so decreases the opportunity cost of leisure. Therefore, unlike an income tax, an endowment tax avoids creating an incentive to substitute leisure for labour. In addition, an endowment tax has an income effect because it necessitates more work for the same level of income. Productivity is incentivized where it matters most, since the highest tax liability is assigned to individuals with the greatest ability to produce. These efficiency considerations drive the economists' acceptance of "income-earning potential" as the first-best tax base and actual earnings as a second-best proxy, merely "the most reliable indicator" for that potential (Mirrlees, 1971, p. 175).

The autonomy objection is that an endowment tax excessively constrains occupational freedom, the so-called "slavery of the talented" (Dworkin, 2002, p. 90). This criticism imagines individuals like Michael Jordan whose rare combination of marketable talents might force them into basketball regardless of their preferences (Fried, 1995). Given feasible alternatives, such as the income tax, this threat to autonomy is rejected as unjust (Murphy & Nagel, 2002, p. 123). However, some commentators are unconvinced that an endowment tax presents any distinct threat to autonomy. Kristi Olson argues that which tax "interferes more with the individual's range of choices is, in fact, empirically contingent" (2010, p. 243). And Louis Kaplow concludes that by using appropriate wage multipliers (avoiding "enslavement" by including a realization requirement when taxing human capital) "one could move substantially toward an ideal income

tax...without raising serious objections concerning violations of individual liberty” (1994, p. 1512).

Most commentators maintain that feasibility constraints seriously diminish the prospects of an endowment tax, even if the autonomy objection can be overcome (Murphy & Nagel, 2002, p. 21). Accurately measuring human capital raises such significant practical challenges that Barbara Fried calls a tax on endowment one “no sane person would seriously suggest levying” (1995, 243). Gathering the necessary information would probably be impossible. Moreover, attempts to do so would seriously violate privacy and would create an incentive to hide or minimize ability. Supposing the information could be obtained, assessing the tax would still require stigmatizing judgments (Anderson, 1999). Finally, it is unclear when endowment should be measured. Few relevant features of endowment are fixed at conception; most are modified by choice, environment, and luck over the lifespan. But if people can react strategically to the tax, then it loses the efficiency justification that motivates it in the first instance.

### **2.3.2.2 Equal Sacrifice**

Obstacles to measuring and taxing endowment motivate alternative interpretations of ability. Most notable is John Stuart Mill’s classic statement of ability to pay as equal sacrifice (2018, Book V), which was later developed by Arthur Pigou, Hugh Dalton, Francis Edgeworth, Nicholas Kaldor, and others (Musgrave, 1959). Mill argued that the burden of paying for the state should inflict the same loss of welfare on each person. Welfare as an index for ability is notoriously difficult to measure, and Adam Smith had earlier claimed that ability should be measured by income (2017, V.II.2.I). Thus, the distribution of income was soon accepted as the best proxy for assessing equality of sacrifice in welfare (Musgrave, 1959, p. 94).

Mill also followed Smith in adopting the notion of “taxable income...as ‘clear’ income, or as income above subsistence” (Seligman, 1908, p. 151). He asserted that subsistence goods are incommensurable with luxuries. Given this distinction, he favoured a basic tax credit that exempted income sufficient to afford “the requisites of life and health” (Mill, 2018, Book V.3). Securing this minimum level of untaxed resources requires protection from the system of tax rules taken together, so he also recommended prohibiting consumption taxes on necessary goods.

The logic of equal sacrifice relies on the idea that each dollar is worth less to the rich than to the poor. Therefore, imposing an equivalent welfare burden on all requires a higher tax burden on the rich. This conclusion rests on a more detailed account of the relationship between welfare and income, and in particular the declining welfare value of income at the margin. But even without the details of such an account, it is plausible that the rich can transfer significant resources to the state with a loss of welfare that is relatively minor compared with the loss experienced by the poor transferring the same amount of resources. And so equality of sacrifice remains an influential consideration in favour of vertical equity and progressive taxation (Slemrod & Bakija, 2017, p. 94).

The initial plausibility of the equal sacrifice account of fairness is limited by two objections. The first is that it ignores the importance of expenditures for tax justice. It prescribes a principle for distributing costs while “the expenditure side of the budget is taken as given” (Musgrave, 1959, p. 90). But a fair distribution of costs in isolation can only be an illusion. That illusion can in some measure be imputed to Mill, since it was he “who separated the analysis of tax equity from the expenditure side of the budget, a separation which for better or for worse (and mostly worse) has dominated tax analysis ever since” (Musgrave, 1990, pp. 114-115). This separation, what James Buchanan called the “fundamental fiscal asymmetry” (Tuerck, 1967, p.

27), attracts a credible charge of myopia, since it aims to defend a narrow principle of tax justice without regard for the overall effect of social institutions.

The second objection is that equal sacrifice assumes a presumptively just pretax allocation of income against which reductions in income can be measured. It frames taxation as an interference with a given distribution without recognizing the necessity of taxation in constructing that very distribution. That framing disregards the extensive role of the state in facilitating the conditions through which large scale productive activity is possible in the first place. This objection maintains that the state is a property-generating entity that constructs the allocation wrongly assumed by equal sacrifice to enjoy independent existence and normative validity.

In general, the traditional public finance approach overlooks both the constructive role of the state and the embedded character of taxation. Equal sacrifice in particular is the most influential interpretation of ability to pay, and it is subject to both criticisms. That is, it takes for granted some pretax distribution of property, and it then proposes a narrow principle of tax fairness to allocate the tax burden across that distribution. But as noted there is an alternative public finance account of how the tax burden should be distributed: the theory of benefits. One promising feature of the theory of benefits is that it takes expenditures seriously, and even as central to tax justice.

### **2.3.3 The Benefit Principle**

The benefit principle adopts an exchange model of the citizen and the state (Musgrave, 1959, p. 62). The state protects rights to life and property, and it provides various goods that benefit its citizens. In return, citizens pay taxes. Since the benefits provided by the state are

unevenly distributed, the benefit principle instructs that the costs of securing them should be assigned in proportion to value received (Seligman, 1908, p. 150). Thus, taxes are understood on the model of *quid pro quo* or insurance premia. This general way of thinking about taxes can again be traced to the middle ages and was endorsed in Adam Smith's first maxim that individuals' taxes should be "in proportion to the revenue they enjoy under protection of the state" (2017, V.II.2.I).<sup>43</sup> Edwin Seligman (1908) finds the view in many writers including Grotius, Hobbes, and Rousseau; it was later developed by Knut Wicksell, Erik Lindahl, and others (Musgrave, 1959).

The idea of "taxation as a price for services rendered" (Musgrave, 1959, p. 63) depends on a comparative baseline against which services rendered (benefits) can be assessed. In that baseline condition no taxes are owed, and no state benefits are received. So it is the no-tax and no-state world: the state of nature. Absent a state to enforce property rights, resources are few, welfare is low, and the distribution of both is roughly equal (Murphy & Nagel, 2002, p. 18). This low but roughly equal baseline suggests that the benefit of taxation to any individual is roughly their actual level of resources or welfare. Thus, the rich are substantial beneficiaries and so on this view owe a great deal more tax than the poor (Seligman, 1908, p. 150). How much more is a matter of dispute, though many have concluded that if taxes should scale proportionately with benefits a proportional or flat tax is implied (Hayek, 1960, p. 316; Musgrave, 1959, p. 62). However, the declining marginal utility of resources implies that a proportional contribution measured by welfare entails progressive taxation, just as it did for the equal sacrifice principle (Murphy & Nagel, 2002, p. 193).

---

<sup>43</sup> Smith's view is ambiguous, since in the same maxim he also says taxes should be "in proportion to their respective abilities." Musgrave (1990, p. 114) calls this an "uneasy mix" of the two principles. Later Smith more clearly endorsed the benefit principle (Musgrave, 1959, p. 66).

The idea that price paid (taxes) should be proportionate to services rendered (benefits) is in itself silent about the quantity and character of services. In addition, the value of many fundamental state benefits is indeterminate, as with domestic order and national defense. This raises a challenge for an exchange model of the state: which benefits should be provided? Given that the benefit principle assumes a model of exchange, one minimal standard might be strict-paretianism. If the right model of the state is a market model of *quid pro quo*, then the exchange ought to be voluntarily in principle, requiring at least that it makes all citizens better off (Fried, 1999). But this raises a question of aggregation: must each benefit (e.g., each road or other public good) satisfy the paretian standard, or only all benefits taken together? If the former, very little can be justified. A road through Kansas is unlikely to attract universal consent. If the latter, nearly any state configuration, including nearly any level of redistribution, can be justified (Fried, 1999, p. 163). Given the state of nature alternative, voluntary exchange imposes a minimal constraint.

In practice, benefits are not measured against “some hypothetical Crusoeian state of nature” (Fried, 1999, p. 163). And there is no market mechanism for determining the correct level of benefits (Musgrave, 1959, p. 66). Instead, the benefit principle is typically applied to existing patterns of expenditure or specific proposals, as though the tax burden can be fairly distributed relative *to those expenditures* without regard for whether that level and character of expenditures is itself justifiable. But this way of proceeding is problematic because an account of justice in taxation cannot take for granted “the nature and extent of government services” (Murphy & Nagel, 2002, p. 15). There is no fair division of costs without a fair distribution of benefits. That is, even if the distribution of costs is a meaningful idea, that distribution depends on an antecedently fair distribution of benefits.

Finally, the benefit principle appears inconsistent with any theory of justice that requires support for the poor (Slemrod & Bakija, 2017, p. 92). This is because the benefit principle characterizes taxes as fees for services rendered, but often state services are rendered precisely to those who cannot afford them. For example, provision of income assistance cannot be coherently paired with a demand for payment. The state cannot issue support payments with one hand and a bill for services with the other. In practice, proponents of the benefit principle realize that, by design, many state benefits cannot command a user fee. Thus, they nearly always include an exemption for some basic level of resources (Fried, 1999). But they provide no “principled basis for absorbing redistributive motives into benefits theory to the extent of basic needs, *but no further...*” (Fried, 1999, p. 162). In the absence of such an account the benefit principle is incompatible with principled support for the poor.

#### **2.3.4 Tax Constitutionalism**

In parallel with public finance, the constitutional constraint approach to taxation developed by James Buchanan and Geoffrey Brennan assumes a set of pretax entitlements against which citizens “suffer the burden of taxation” (2000, p. 3). But it is notably distinct from traditional tax burden analysis, and from the theory of optimal taxation. It begins by rejecting political models that represent government as a benevolent despot (Hamlin, 2018). It identifies the untenable assumption of government benevolence in accounts of tax equity that proceed as though government planners will act in socially optimal ways rather than privileging their own interests. Brennan and Buchanan argue that the benevolence assumption reflects an unjustified motivational asymmetry between political and market agents. For instance, in optimal tax theory market agents are modeled as self-interested maximizers, and there is no obvious reason to

assume government agents are fundamentally different in their motivations. The implausible assumption of government benevolence then tends to be paired with a view of government power as unconstrained. But even if a social planner knows the theoretically optimal tax design, and is motivated to implement it, political constraints might render the project infeasible. Crucial among such constraints are norms of democratic decision making which seriously limit the range of feasible rules of taxation. Brennan and Buchanan conclude that proceeding as though government has both the power and motivation to unilaterally implement some optimal set of tax rules is misguided from the outset.

Instead, the tax constitution approach contributes an alternative framework for developing an account of tax justice (Brennan & Buchanan, 2000). It posits that fair tax rules must be established from a particular point of view. That view, called the constitutional orientation, is characterized as the perspective of a rational agent choosing rules from behind a partial veil of ignorance. The nature of the agent's choice is clarified first by distinguishing between rules of a game and plays within that game. Rules provide a framework or a set of boundaries within which legitimate "in-period" moves can be identified (Brennan & Buchanan, 2000, p. 3). From the constitutional orientation, the agent evaluates various general rules of taxation, combinations of which set the background for economic strategies or plays once those rules are fixed. The chosen rules are argued to be fair because the agent in the constitutional choice setting is behind a translucent veil of ignorance. That is, the agent is ignorant of arbitrary personal characteristics but is rational and informed about the nature of government in the postconstitutional (in-period) sequence.

The description of government provided in the constitutional choice setting is pivotal for the choosing agent. To illustrate this point, Brennan and Buchanan adopt a model of government

analogous to the standard market model of monopoly: the revenue maximizing leviathan (2000, p. 33). The leviathan is constitutionally constrained but otherwise aims to extract as much tax as possible. They acknowledge that there exists a range of motivational descriptions, and that the leviathan model is an oversimplification, but nevertheless they examine its implications for insights into tax equity. Given their exploitative model of government, they observe that rational agents have an interest in constitutional constraints that limit that exploitation. Therefore, agents would endorse tax rules that limit the taxing power in a number of ways. For instance, rules might circumscribe the tax base since an expansive tax base provides government with more opportunity to exploit taxpayers. And common tax rates might be required since variable rate structures are yet another exploitable source of government revenue. Or the “tax take” could be limited from the outset not to exceed a set portion of overall economic activity. Brennan and Buchanan also argue that the revenue maximizing leviathan model of government can illuminate other aspects of social organization. For instance, they note that a federation (rejecting the assumption of government unity) can enhance tax competition among its regions, thereby reducing the monopoly control of the state on the power to tax. Of course, as they acknowledge, they do not provide a realistic model of government. Therefore, the implications of their model are correspondingly limited. However, their structure of analysis calls for consideration of political realities in developing a theory of fair taxation, a project most comprehensively developed by the theory of public choice (Mueller, 2003).

## **2.4 Tax as Social Architecture**

The social architecture approach to taxation rejects the basic premise of tax burden analysis, namely that taxation diminishes the presumptively legitimate pretax holdings of those

on whom taxes are imposed. It therefore abandons the project of identifying narrow principles of fairness for application to pretax holdings, maintaining that tax justice is “not a matter of applying some equitable-seeming function to a morally arbitrary initial distribution...” (Murphy & Nagel, 2002, p. 30). Rather, tax justice must be assessed in terms of outcomes: post-tax allocations. This approach emphasizes that pretax resource allocations and expenditures cannot be taken for granted as arising independently from a system of taxation. Taxation is already presupposed in the construction of ostensibly pretax distributions. And the tax system cannot be conceived of as normatively detached, as though the consequences of taxation might be evaluated independent of the social institutions in which they are embedded.

#### **2.4.1 Constructing Entitlements**

One diagnosis of the problem of tax burden analysis is that it implicitly embraces a naïve view of economic transactions in a modern economy. Thomas Nagel and Liam Murphy call this view unreflective or “everyday” libertarianism and characterize it as assuming that societies provide free markets in which people can exchange goods and services for mutual benefit. Everyday libertarianism assumes that this process of exchange secures at least presumptive justice in distributive outcomes (Murphy & Nagel, 2002, p. 15). If someone agrees to pay some amount for a good or service then, in the absence of special justification, the payee is *entitled* to receive that amount. The market therefore supplies a normatively significant pretax distribution of resources. On this naïve view, taxation is an after-the-fact interference with those entitlements and so requires special justification.

In contrast, the social architecture approach rejects the everyday libertarian account of entitlements as logically incoherent. This is because there is no space, even conceptual, wherein

the free market operates and then taxes are imposed on the market distribution after the fact. The state—and so taxation—are preconditions of any modern market distribution: there is “no market without government and no government without taxes” (Murphy & Nagel, 2002, p. 32). As described in chapter one, the terms of the market, especially property and contract law, are set by the state. As a result, “no sharp line can be drawn between markets and government: the two entities have no existence detached from one another” (Holmes & Sunstein, 2000, p. 69). Or, as Friedrich Hayek put it, “the question whether the state should or should not ‘act’ or ‘interfere’ poses an altogether false alternative, and the term ‘laissez-faire’ is a highly ambiguous and misleading description of the principles on which a liberal policy is based” (2014, p. 118).

Given the inextricable relationship between government and market outcomes, it is a misleading exercise to imagine a modern society but without taxation (Murphy & Nagel, 2002, p. 16). Hypothetical models that assume a tax-free world do just that. Carl Shoup described the resulting confusion: “To say, for example, that households with before-tax incomes between \$2,000 and \$5,000 pay 12 percent of that income in taxes, directly and indirectly, is to make a statement that is without significance because it is conceptually invalid. It is conceptually invalid because it postulates, for implicit comparison, a state of affairs in which there are no taxes whatever, and no government borrowing or creation of new money, hence impliedly no government services, not even of the minimum type and amount necessary to assure existence of the society . . . . [This] objection is conclusive” (1969, p. 577-578). Shoup objected to the implicit postulation of a pretax baseline because a modern society without taxation is a contradiction in terms. The state develops and enforces property, tort, and contract law. It regulates banks, defends its territory and citizens, and supplies public infrastructure for power, water, and the internet. States facilitate global advertising and the transportation of goods

between nations. It is against that background that any ostensibly pretax distribution emerges. Thus, the no-tax world is the rather grim no-state world: it is simply the state of nature.

Seen in the preceding way, pretax entitlements are pre-state entitlements. But any pre-state distribution of resources is both epistemically inaccessible and normatively irrelevant. There is of course no sense in trying to determine the equilibrium price of a piece of digital music in the state of nature. The inquiry is incoherent because the existence of digital music assumes a set of institutions that includes taxation. Thus taxation is essential to constructing the ostensibly pretax baseline and cannot be understood as an interference after the fact. This undermines the everyday libertarian view that individuals are presumptively entitled to all of the resources they possess prior to taxation.

The everyday libertarian approach can also be critiqued as a conflation of possession and entitlement. A business owner mistaking her business revenue for an entitlement to income would commit a parallel error. Though she may possess significant revenue, that revenue is only an accounting figure, and her entitlement computation will render a different figure. She must subtract expenses due to her suppliers, staff, landlord, and so on. Those transfers are not interferences with her normatively significant pre-expense revenue. Indeed, her revenue alone can convey nothing at all about her entitlements. In the case of a business, expenses are conditions that must be satisfied to generate entitlement to any income at all. Similarly, taxes are conditions antecedent to the generation of resource entitlements, not interferences with a pre-existing and normatively significant allocation of those entitlements.

Absent a meaningful pretax distribution, burden analysis as a general approach to tax equity appears misguided, along with its principles of horizontal and vertical equity. The structure of burden analysis involves apportioning the tax burden over some given distribution of

pretax resources, welfare, or other index. And its equity principles “if justifiable, must derive from the conception that a person is morally entitled to his or her pretax level of well-being, and that the state, through its tax structure, must respect that position” (Elkins, 2006, p. 87). If the pretax distribution of the tax base is normatively irrelevant (or incoherent), then the difference between that amount and the post-tax allocation cannot be the correct measure for tax equity.

Many familiar claims of tax justice appeal to horizontal equity and are therefore vulnerable to criticism of the pretax baseline. For example, it is often claimed that a consumption tax is fairer than an income tax (Bradford, 1986, p. 179). This is distinct from the efficiency claim that a tax on capital income reduces the return to saving and so increases the incentive to substitute consumption for saving. The fairness claim, in contrast, is that a tax on capital income burdens equivalently situated individuals differently. And the equality of their situation is assessed in terms of their pretax holdings. A saver and a consumer with equal pretax income only differ on the arbitrary basis of their preference for postponing consumption. Both are taxed on initial income, and then the saver is taxed on capital income, and so the income tax base is “clearly unjust” (Fisher, 2007, p. 253). But if the fact that the saver and the consumer possess equal pretax income is of no normative significance, then undermining that equality by the choice of tax base does not support the charge of injustice.

Rejecting entitlements to pretax holdings as mythical has drawn criticism. Some is based on the misperception that such rejection implies “that the entire social product really belongs to the government, and that all after-tax income should be seen as a kind of dole that each of us receives from the government, if it chooses to look on us with favor” (Murphy & Nagel, 2002, p. 176). But as Murphy (2005, p. 123) explains, “economic and legal institutions as a system will be just depending on how well they secure political values, such as liberty, welfare, opportunity,

personal responsibility, and equality - both of status, by which I mean primarily sexual and racial equality, and economic. Different political theories interpret and weigh these values differently, so abandoning the traditional way of thinking about tax justice leaves open a full range of views, from left to right, on what a just overall economic system, taxation included, will look like. In particular, it leaves open what the rate structure of the best tax scheme will be.” This is a clear rejection of the view that the government is permitted to do what it pleases with the social product.

A more substantive line of criticism coming from Geoffrey Brennan and others is that the social architecture approach considers only “final distributions” (2005, p. 207) and ignores “the process by which the final distribution emerges” (2005, p. 216). Marc Fleurbaey calls this the “Mirrlees-Murphy-Nagel approach” (2018, p. 37) and notes that the Mirrlees model “focuses on utility and completely disregards the way in which utility is generated” (2018, p. 38). Brennan characterizes this approach in terms of “the principle of ‘anonymity’—the claim that permutations of position [in the resource distribution] are normatively irrelevant” (2018, p. 79). However, rejecting tax burden analysis does not entail neglecting considerations of process. Indeed, Murphy notes that “the weight of values such as liberty and responsibility are appropriately part of the discussion” (2005, p. 149). And Murphy and Nagel write that justifications for the “final distribution” should appeal directly to process (2002, pp. 128-129), including “considerations of individual liberty, desert, and responsibility as well as to general welfare, equality of opportunity, and so forth” (2002, p. 75). Principles of justice that are choice-sensitive, for example, would compel a post-tax allocation that tracked in some ways the choices of taxpayers. That result is compatible with rejecting the normative significance of any pretax baseline, but inconsistent with the principle of anonymity as Brennan describes it. All that is

ruled out by the constructivist critique, then, is appeal to the presumptive justice of any pretax distributive baseline.

The constructivist account of entitlements is also criticized for ignoring prepolitical property rights. The notion that property rights are generated by the state seems to trample John Locke's claim that "every man has a property in his own person" (1980, V.27) and the derivative state-independent capacity to create property rights in the world. This, it is argued, is an unjustified adoption of the view of David Hume and Jeremy Bentham that all property is a creature of law. As Bentham put it, "property and law are born together and die together. Before the laws there was no property; take away the laws, all property ceases" (2012, p. 113). Murphy and Nagel do evidently take that position, insisting on "the conventionality of property" (2002, p. 175), the "denial that property rights are morally fundamental" (2002, p. 175), and the rejection of "moral property rights to which the legal rules of property ought to conform" (2005, 148). Instead, they claim property "is created by the tax system" (2002, p. 175), "purely conventional" (2002, p. 175) and "individual citizens don't own anything except through laws" (2002, p. 176).

As I explain in chapter three, I think Murphy and Nagel conflate the existence of conventional property rights with the nonexistence of moral property rights. They rightly observe that there are legal property rights, the majority of which could not exist without "money, banks, corporations, stock exchanges, patents, or a modern market economy" (Murphy & Nagel, 2002, p. 32). And no one could doubt the tautology that legal property rights are a consequence of the state and its laws. But this observation alone does not undermine the existence of "moral property rights to which the legal rules of property ought to conform" (Murphy, 2005, p. 148). As Brennan (2005) illustratively points out, the existence of state-dependent legal rights that protect life do not in themselves undermine the prepolitical normative foundations of the right to

life. Property rights are not obviously contingent on law, as there are “notions of property [that] pre-date the establishment of modern systems of law and enforcement” (2005, p. 211). Indeed, moral property rights could exist even if “the *existing* set of [legal] property rights is ethically arbitrary and has emerged ‘purely conventionally’” (Brennan, 2018, p. 70). However, the central thesis of the social architecture approach appears untouched by the status of property rights since “the argument, for societies in which there is a state, is that without the state, and therefore taxation, we wouldn't have the pretax incomes we do, and so it makes no sense at all to use pretax income as the baseline for thinking about just taxation” (Murphy, 2005, p. 148). The analysis of property in chapter one supports Murphy’s claim about the role of the state in constructing pretax income. Thus, even moral property rights cannot render the pretax baseline coherent.

As described in chapter one, ownership is generally analyzed as a collection of specific entitlements, such as entitlements to possess, use, exchange, give, destroy, and so on (Honoré, 1961). A theory of distributive justice bears directly on how those entitlements are allocated. Certainly the questions raised by distributive justice are not settled by referring to “the way in which any particular legal system allocates claims” (Murphy, 2005, p. 149). So even Murphy and Nagel accept that a correct account of the various entitlements to resources is a moral rather than legal one, and indeed that the law can violate that account. As I argue in chapter three, their defense of a thoroughly moral account of resource entitlements implies that their conventionalism is much less arbitrary than their critics suggest.

Given the constructivist ambiguity about the status of property rights, theorists with libertarian sympathies have tried to preserve the sanctity of pretax market distributions by arguing that if one “fully owns each factor of production” then one fully owns the product of

those factors (Vallentyne, 2018, p. 102). On this view, market exchanges of fully owned products must be protected from taxation. In outlining this argument, Peter Vallentyne sets aside cases where the factors of production are not fully owned as unimportant for tax justice (2018, p. 102). But such cases are central to tax justice if the state is a factor of production, a “‘silent partner’ in every business enterprise” (James, 1886, p. 486). Taxes are a necessary condition of all the elements of the state that make immense productivity possible. Edmund James suggested that “to test the relative productivity of the state and the individual, compare the fortune accumulated by Cornelius Vanderbilt in America with what he might have accumulated had he been adopted when an infant by a family [living in an underdeveloped state]” (as cited in Fried, 1995, p. 237). Wherever the state is a factor of production, no individual “fully owns each factor of production.”

#### **2.4.2 Myopia and Holism**

The social architecture view pairs its constructivist critique of the pretax baseline with an objection to judgments of fairness that are bounded to the domain of taxation. Bounded judgments are myopic because they overlook factors relevant to assessing fair treatment by the state. The claim is that taxes are neither fair nor unfair in themselves, but rather one part of a comprehensive system that must, as a whole, treat people fairly. As Alan Hamlin puts it, “the tax and benefit system cannot properly be seen in isolation from other aspects of the institutional structure of society” (2018, p. 33).

Myopic judgments of tax fairness generally come in two varieties. The first evaluates specific tax burdens instead of the net effect of the system of tax and transfer. But the effect of any tax rule can be offset by some other rule. For example, a mortgage interest deduction might

be paired with a high property tax. Considered in isolation, the fairness of that deduction is indeterminate. Moreover, offsetting can be even more direct. The system of tax rules also specifies transfers and expenditures, and a compensated burden is not a “burden” in the relevant sense: “economists have long recognized the essential equivalence between taxes and transfer payments” (Bradford, 1995, as cited in Murphy & Nagel, 2002, p. 14). A tax burden paired with an equal transfer is the clearest case of a tax proposal that has no net effect, and net effect is the relevant measure.

The second variety of myopia does consider the full interconnected system of taxes and transfers, but it does so as though it were detached from the rest of the institutional structure. Arthur Pigou noted the problem with such detached judgments of fairness: “People’s economic well-being depends on the whole system of law, including the laws of property, contract and bequest, and not merely upon the law about taxes. To hold that the law about taxes ought to affect different people’s satisfactions equally, while allowing that the rest of the legal system may properly affect them very unequally, seems not a little arbitrary” (Pigou, 1947, as cited in Murphy & Nagel, 2002, p. 26). Further, the force of Pigou’s observation extends beyond the legal system to include other functions of the state.

Despite these considerations, myopia is the norm in both the theory and practice of tax justice. Fairness objections to “double taxation” are a case in point (Musgrave, 1959, p. 163). There can be nothing unjust about multiple taxes on some good unless those taxes sum to an unjust rate of taxation. Simply put, whether the appropriate amount of tax is collected across one or more taxes is of no normative significance (Musgrave, 1959, p. 163). The fairness of rate structures is similarly derivative. There are no grounds to “fix on rate structures themselves as

fair or unfair” since they “are just a means to operationalize other prior, moral commitments about the proper role of government” (Fried, 1999, p. 158).

On the practical side, many governments are guilty of myopia, since they often provide “a wealth of data...on the distribution of tax burdens” (Zelenak, 2003, p. 2269). In the United States, proposed legislation is accompanied by distribution tables that claim to show how the resulting tax burden will be allocated. And “because tax legislation is, in general, distinct from expenditure legislation, tax policy inevitably attracts judgments of fairness or unfairness in its own right” (Murphy & Nagel, 2002, p. 162). Thus the predicted distribution of burdens is assessed for fairness without regard even for directly offsetting transfers including Social Security and Medicare.

Avoiding the myopia objection requires theorists to resist evaluating the fairness of specific taxes in isolation. Tax fairness is not like respect which can be looked for in every case. Myopic tax assessments result in misguided judgments of fairness that are not robust, as they overlook the net effect of social institutions. These wider distributive implications are essential to judgments of tax fairness and so this argument maintains that justice in taxation is inextricable from the cumulative impact of the comprehensive institutional structure.

### **2.4.3 Optimal Tax**

Some contemporary economic analyses of tax equity join the social architecture approach in setting aside the traditional assumptions of public finance (Hamlin, 2018, p. 17, 20-26; Fleurbaey, 2018, 37). The result is abandonment of the general method of burden analysis in favour of what is now the mainstream academic economic approach to evaluating taxation. That approach is “optimal tax” which aims to systematize Jean-Baptiste Colbert’s quip that “the art of

taxation consists in so plucking the goose as to procure the largest quantity of feathers with the least possible amount of hissing” (as cited in McKechnie, 1896, p, 77).

At a general level, optimal tax contributes a structure of analysis (Hamlin, 2018). It provides a model environment in which to predict the effect of various tax rules on a specified social welfare function. Models are structured as a double optimization problem, where the solution optimizes satisfaction of the social welfare function given that individuals will optimize their own value functions within the ruleset adopted by the model (Fleurbaey, 2018, pp. 40-41).

Optimal tax models ascribe no special significance to any pretax baseline. Instead, the normative force of the model is found in the social welfare function. Ideally that function ranks possible social states from less to more desirable. Rankings can (and frequently do) emphasize aggregate utility, but they need not, and many include a measure of inequality aversion or sensitivity to other distributive considerations. Models partly avoid myopia, at least to the extent that they consider both taxes and transfers through inclusion of negative tax rates. With all rates on the table, agent responses to the selected set of rules will change the overall effectiveness of those rules at achieving whatever normative ideal is adopted by the model.

In order to formulate mathematically tractable problems, optimal tax models make a variety of simplifying assumptions. Assumptions differ depending on the exact account, but they might include a maximizing conception of practical rationality, homogenous agent preferences, identical utility functions, and so on. The limitations of such assumptions are well-known, and they raise a general concern about how well model-derived conclusions track the world in which the assumptions do not hold.

There is also reason for caution when deriving model optima for guidance in a suboptimal world. Marc Fleurbaey calls this the “danger in finding a general formula” (2018, p.

48). A familiar example of this problem comes from general equilibrium models. Those models specify a series of optimality conditions necessary for market equilibria and demonstrate that all equilibria are Pareto efficient. But when all of the conditions do not hold, closer approximation of those conditions does not predict closer approximation of Pareto efficiency. Instead, given imperfect satisfaction of any optimality conditions, efficiency improvements may require further violations of those conditions (Lipsey & Lancaster, 1956).

In the optimal tax case, a parallel confusion must be avoided. For example, the optimal formula might include decreasing marginal tax rates. The temptation, then, is to conclude that “any reform that moves in the direction of decreasing marginal tax rates would improve the situation. But, of course, nothing can be less true” (Fleurbaey, 2018, p. 48). In general, the optimum and its conditions do not provide a reliable guide for approximation (Estlund, 2019). Martin Feldstein explains that “everything we know about the theory of economic policy in other areas reminds us that optimal piecemeal policies cannot be made by haphazard steps in the direction of the global optimum, that a constrained second-best policy cannot be guided by the conclusions of an unconstrained optimization” (1976, p. 77). Such haphazard steps may be avoided when a social welfare function provides a ranking of outcomes, rather than only an optimal solution. A ranking allows assessment to proceed comparatively rather than transcendently, evaluating each predicted outcome relative to alternatives, rather than “relying on a vague comparison with a stylized fact about the optimum” (Fleurbaey, 2018, p. 48).

#### **2.4.4 Choice Architecture**

Taxes do much more than raise revenue. They also change incentives, and so they shape behaviour. As a result, they are an important factor in choice architecture (Thaler & Sunstein,

2021). Evidence suggests that in some cases individuals (and so markets) respond dynamically to these incentives. For example, raising the income tax rate changes the labour supply: some people will work less, and some not at all (Kaygusuz, 2010; Saez et al., 2012). Some people restructure their finances to take advantage of comparatively low tax rates on capital gains. And there are plenty of historical examples that highlight this dynamic character of markets, some of which can be observed in creative responses to taxes on building height and width, number of windows, construction brick size, and so on.

Since taxes change the quantity and character of economic behaviour, this again raises the question of how the state and the market are related. The standard view is that the state and the economy are somehow dissociable. This view is reflected in the widely held belief that taxes “unquestionably impose costs on the economy” (Slemrod & Bakija, 2017, p. 227). These costs are missed opportunities for mutually beneficial exchange, also called deadweight loss. When taxes must be levied, the state should aim at “neutral tax treatment of economic activities” (Slemrod & Bakija, 2017, p. 228). The main exception to this general rule is the over or under production of goods due to market failure, for which taxes are readily prescribed. But the concepts of deadweight loss, market failure, and neutrality all assume some notion of a tax-free market equilibrium.

Deadweight loss is a measure of lost economic efficiency, and taxation is considered a paradigm cause of such loss. The logic seems straightforward enough. First, assume a tax-free market equilibrium for some good. Then impose a tax, shifting the supply curve up to reflect the additional cost of the good. As price rises the quantity consumed falls. Consequently, the tax not only converts some of the producer and consumer surplus to state revenue, it also reduces the quantity of transactions. For example, if a sales tax on a content streaming service prevents a

purchase, then the supplier, consumer, and government all gain nothing: the missed transaction is simply deadweight loss.

Over and under production are similarly understood with reference to a tax-free equilibrium. For example, some transactions impose costs on nonconsenting third parties. A classic example is alcohol, which attracts social costs related to healthcare, crime, motor vehicle accidents, and so on. These costs are “divergences between private and social net product” (Pigou, 2016, II.IX.13). The private net product of alcohol sales is “unduly large” relative to the lower social net product. Additional costs are externalized, and so alcohol sales are effectively subsidized, resulting in overproduction of alcohol. Higher alcohol taxes are then recommended to reduce the quantity produced, bringing the market in alcohol back into equilibrium.

As I argue in chapter four, these sorts of analyses can be misleading from the view that taxation and markets emerge together. The tax-free world is a stateless one, and so the tax-free equilibrium assumes a world without government. It is meaningless, for example, to investigate the equilibrium price of Random Length Lumber futures in the state of nature. Taxation is a precondition of any nonzero market equilibrium in such goods. But if taxation makes any quantity of transactions possible, then taxes cannot be measured as a distortion or corrective relative to a tax-free baseline. Just as the social architecture approach rejects a pretax distribution of resources which taxation must preserve, so too it must reject a tax-free account of the quantity and character of economic activity. Both are products of taxation, not its normatively valid antecedents. The behavioural consequences of the transactional scaffolding provided by taxation, then, must be evaluated on other grounds.

A search for other grounds suggests a wider understanding of market failure. Taxes predictably modify behaviour, and goods can be over or under produced relative to a variety of

normative standards (Brennan, 2008). This might include familiar standards such as “(a) securing basic capabilities for all persons (Sen/Nussbaum), (b) providing equal opportunity for welfare (Arneson), (c) maximizing aggregate human welfare (Mill), (d) distributing economic resources so as to provide the greatest benefit to the least advantaged (Rawls’ difference principle) or (e) establishing and maintaining equality of resources (Dworkin)” (Macleod, 2008, p. 2). Against the backdrop of broader conceptions of distributive justice, the success or failure of markets can be measured by the degree to which they satisfy the appropriate normative standards.

Here as elsewhere the assessment must proceed holistically, and this requires considering the net effect of the tax, including the disposition of revenue. For example, perhaps there is good reason to increase access to streaming services. Considered in isolation, a high sales tax on those services might be predicted to reduce demand. But if generated revenue is invested in enforcing strong antitrust regulations that dismantle a monopoly, then the tax might increase demand, all things considered. So there is no avoiding the fact that taxes will change the quantity of economic transactions, but the net effect depends on a wide set of considerations. Whether those changes are attractive depends on the application of acceptable normative standards.

Many otherwise familiar tax debates look different from this perspective. For example, high taxes on cigarettes and sugary drinks are often criticized as regressive because they disproportionately affect low-income families (Sanders, 2016). People with low incomes consume more of these products and, in addition, flat commodities taxes cost proportionately more for those with fewer resources. A parallel criticism is leveled against tax subsidies. Governments often provide tax benefits for consumers of energy efficient windows, appliances, or vehicles. Those benefits accrue to those who need them least, since the poor cannot afford the subsidized items. But according to the social architecture approach, little can be said about these

regressive taxes considered in isolation. A sugary drink tax might be one element of a progressive revenue recycling project (Lockwood, 2017). Perhaps counterbalancing expenditures increase the quality of life in low-income neighbourhoods, and in addition lower healthcare costs and insurance premiums might result in a net benefit for the poor. And incentivizing energy efficiency with tax subsidies might be justified by important environmental considerations that outweigh the policy's distributive consequences. Indeed, people with low incomes are often most vulnerable to climate change and so might stand to benefit significantly from such subsidies, even when those subsidies are largely appropriated by individuals with high incomes. Evaluating these proposals against "authoritative and overarching normative standards" (MacLeod, 2008, p. 1) inevitably raises concerns about paternalism and the appropriate scope of state efforts to change behaviour. But such consideration cannot assume that taxes are raised or lowered relative to a presumptively justified tax-free baseline.

#### **2.4.4.1 Public Goods**

The funding of public goods is widely considered a core function of taxation. Economists typically identify public goods by two primary characteristics. First, public goods tend towards non-excludability, meaning their consumption cannot be made conditional, or it is difficult to do so. Thus they "cannot be provided to anybody unless they are provided to everybody" (Murphy & Nagel, 2002, p. 46). And second, they are non-rivalrous, meaning they tend to be undiminished by consumption and their marginal production costs approach zero. The most uncontroversial examples of public goods understood in this way are national defense and domestic order. Both are in everyone's self-interest collectively, but informational constraints and the possibility of free riding predict they will be undersupplied in the absence of intervention by the coercive power of the state (Halliday, 2013).

Even the fact that an increased supply of some good is in everyone's self-interest does not obviously justify the state exercising its coercive power to supply it (Halliday, 2013, p. 1113). Further argument is required, and this suggests a need for a "moralised account of which public goods warrant coercive supply and which don't" (Halliday, 2013, p. 1112). Such an account would clarify the "level of funds [that] can justly be extracted" (Halliday, 2013, p. 1112). That is, according to Halliday, it would identify which goods people should be forced to pay for, and how much they should be forced to pay. However, relying on the economic definition of a public good in developing this sort of account presents two significant problems.

The first problem with the economic analysis of public goods is its reliance on determining the marginal value individuals place on various levels of those goods. But the marginal value of some good for a consumer depends on an existing resource distribution. Thus the trouble with approaching the problem in the standard way is that it assumes without justification an entitlement to some pretax distribution. That is, the analysis "presupposes an antecedent distribution of bundles of resources, prior to the provision of public goods, and . . . there is no such thing" (Murphy & Nagel, 2002, p. 47). An assessment of the justice of taxing resources to fund public goods cannot depend on a chimerical pre-state allocation of resources with which state action interferes.

The second problem is that the economic account may suggest that the scope of permissible state funding or provision of goods is constrained such that it tracks the "public" characteristics of non-rivalry and non-excludability. This slide from a specific conception of public goods into a normative claim is arguably found in Rawls' *A Theory of Justice* (O'Neill, 2021). However there is no obvious reason why state funded or provided goods should be guided by the notion of a public good in the economists' sense. Indeed, that specific notion may lack

coherence, relying as it does on an antecedent distribution of entitlements. In the absence of such a distribution, the normative analysis of public goods appears to be subsumed in the more general study of choice architecture. A bundle of public goods must be justified alongside the system of taxation as part of a comprehensive institutional structure that constructs a fair post-tax allocation of resources. Predictably, the assessment of taxation for the funding or provision of public goods looks quite different from this architectural point of view.

The architectural point of view leaves open how “public goods” ought to be conceived. Most people agree that a justified configuration of the state and its economy will involve state provision of at least some goods (Murphy & Nagel, 2002, p. 5). But the relevant goods may be understood quite broadly to include “health, education, public safety, national security, criminal justice, environmental protection, recreation, procreation, and other social goods...” (Sandel, 2013, p. 9). This “richer picture” of public goods might appeal to a wide range of principles and political values in delineating the proper scope of state provision and financing (O’Neill, 2021, p. 18; Cordelli, 2020; Dorfman & Harel, 2016). Some goods may contribute to the conditions of authentic choice and citizen autonomy. Others may be at risk of corruption or degradation when subjected to the logic of markets or would otherwise permit an intrusion “of market values, into spheres of life where they don’t belong” (Sandel, 2013, p. 7; Satz, 2010). In addition, the state may be well positioned to preserve a set of valuable opportunities against the caprice of the moment. And unequal access to some goods, especially when access tracks ability to pay, may undermine conditions of democratic equality or other political values. Such arguments might be invoked to support an account of public goods including art galleries, sporting facilities, opera houses, theaters, museums and so on, even when those goods do not satisfy the economic conception of “public.”

The character of public goods in this broad sense is a promising area of inquiry (Hussain, 2018). However such goods are best understood, the social architecture view rejects the standard analysis that rests on individuals' marginal valuations of a good relative to a given distribution of resources. Therefore the state's role in regulating private competition or funding and even supplying certain goods cannot be derived from an imagined tax-free world, as though a distribution of resource entitlements exists prior to a justified provision of public goods in this broad sense.

## **2.5 Emerging Issues**

### **2.5.1 Monetary Theory**

Accounts of tax justice normally take for granted that one central purpose of taxation is to raise revenue (e.g., O'Neill & Orr, 2018, p. 1; Murphy & Nagel, 2002, p. 3). However, recent work in philosophy and macroeconomics has challenged this unified picture, claiming that in some cases taxes do not raise revenue (Hockett & James, 2020; Kelton, 2020). The challenge rests on a distinction between currency issuers and currency users. Currency using governments need money to pay for things, and they acquire that money through borrowing or taxation. But currency issuing governments are fundamentally distinct: they control the money supply.

Robert Hockett and Aaron James point out that, for currency issuing governments, "taxes do not pay for spending" (Hockett & James, 2020, p. 38). Given a high level of monetary sovereignty, governments do not raise money, instead they create it, and so they face no revenue constraint (Hockett & James, 2020, p. 309). They make the case that the state must remain responsive to controlling the overall money supply, but when a state issues its own currency it "doesn't need to 'get money' or 'have money' in order to spend it" (Hockett & James, 2020, p.

38). That misguided idea confuses currency issuers with currency users; it conceives of governments like private households, the so-called “household fallacy.” Regarding government as creating rather than collecting revenue changes the functional role of taxation. Notably, for monetary sovereigns, taxes paid do not contribute to government coffers: governments do not issue taxes in order to spend. Instead, taxpayers relinquish assets in satisfaction of a government-imposed liability, resulting in a zero sum on the overall state balance sheet: debts are simply extinguished, and currency is destroyed (Hockett & James, 2020, pp. 38-39).

On this account, taxation remains crucial for both currency issuing and currency using governments. For currency users (those without monetary sovereignty) taxation does raise revenue as is normally assumed. Like households, currency using governments must raise money in order to spend it. But for currency issuers, taxes are not a source of revenue. Instead, taxes are one instrument for controlling the money supply, helping ensure that money creation does not outpace the productive capacity of the economy. Thus, normative assessment of monetary sovereigns must be responsive to this macroeconomic function in addition to the functions described above. For both issuers and users alike, therefore, taxation remains a critical technology for shaping the distribution of property entitlements and for creating incentives that contribute to a society’s choice architecture.

### **2.5.2 Tax Among Nations**

The value of the power to tax depends on the size of the tax base. In a global economy, the tax base is a diverse and increasingly mobile target. Mobility of capital encourages competition, as states design their regulatory environments, including their tax rules, to increase

their share of the global tax base. Tax competition is generally thought to raise concerns when it threatens state sovereignty or contributes to unjust global inequality.

As with its domestic counterpart, philosophers have mostly avoided normative discussions of international taxation. Peter Dietsch (2015) is an exception, and he develops an account of a just distribution of the tax base between nations. He distinguishes between two competitive tactics, poaching and luring, which depend on the tax base at stake. Portfolio capital and paper profits are subject to poaching since they can be detached from the productive activity that generates them ('virtual' tax competition). Luring foreign direct investment, in contrast, involves using favourable taxation to attract the productive activity itself ('real' tax competition) (van Apeldoorn, 2016).

Tax competition is not always objectionable, but Dietsch argues it has two possible disadvantages. The first is that it can undermine fiscal autonomy. Democratic states should have the capacity to set the size of their budget relative to their domestic product, and to determine their level of redistribution. Downward pressure on tax rates and regressive fiscal arrangements both undermine this autonomy prerogative. The second disadvantage is that tax competition can contribute to global inequality. Capital is more mobile than labour, and so tax competition is responsive to the interests of capital, contributing to an ever-widening resource disparity.

Dietsch's reform proposal includes a unitary tax with formulary apportionment, which would distribute tax base in proportion to various aspects of corporate operations (common alternatives include source-based taxation (e.g., Kane, 2015) and destination-based taxation (e.g., Devereux & Vella, 2014)). Assigning taxing rights to hosts of economic activity is a standard recommendation, and "even those who agree on nothing else when it comes to international taxation would likely agree that each state should be entitled to revenues derived from economic

activity that occurs within its borders” (Dean, 2009, p. 162 as cited in Kern, 2020, pp. 3-4). But Dietsch supplements this recommendation with two regulatory principles. The first is the membership principle, which holds that individuals are only liable to pay tax to states of which they are members. This would prohibit poaching, and in turn focus competition on luring. The second is the fiscal policy constraint which addresses luring by prohibiting fiscal policy that negatively impacts aggregate fiscal self-determination of states whenever that policy is strategically motivated.

Similar to the domestic case, international reform proposals often approach taxation as “a subject of justice in its own right” (Dietsch & Rixen, 2019, p. 503). Dietsch accepts that approach, noting that global distributive justice is controversial, but the value of fiscal autonomy is less so. Therefore, his two principles are “designed to protect the fiscal autonomy of states” (2019, p. 503). They provide guidance to “evaluate the functioning of the international tax system by analyzing whether it serves fiscal autonomy” (2019, p. 503). Dietsch construes fiscal autonomy as a state’s sovereignty over its resources: states have a valuable prerogative to distribute their resources as their citizens decide. But this underestimates the degree to which fiscal autonomy “in part *consists in* having access to sufficient tax revenue” (van Apeldoorn, 2016, p. 491). Dietsch’s more austere notion of fiscal autonomy may be less controversial, but then it offers no guidance for how international tax rules should distribute the tax base among nations. After all, those rules define the antecedent level of resources over which states are sovereign. The relevant distinction is illustrated by the domestic case, where valuing an individual’s control over her resources does not resolve the matter of her correct level of resource entitlements. Just as individual autonomy presupposes a certain level of resources, so too does fiscal autonomy presuppose an adequate level of tax base. How much tax base is

necessary is an open question, but Dietsch's conception of fiscal autonomy cannot resolve that question since it only maintains that state's should remain autonomous over their resources, whatever those resources happen to be. A more robust account of fiscal autonomy could provide guidance for distribution of the tax base, but it unavoidably merges with a more general conception of global distributive justice.

Once again tax justice appears inextricable from distributive justice. International reform is necessary precisely because existing patterns of economic activity and resource possession are arbitrary and unjustified. And here as elsewhere tax justice is "not a matter of applying some equitable-seeming function to a morally arbitrary initial distribution..." (Murphy & Nagel, 2002, p. 30). Therefore, reform proposals grounded in the status quo are an example of the 'myth of ownership' gone global (Risse & Meyer, 2018). Even the near-universal idea that firms should contribute to states that host their economic activity is open to doubt. Perhaps reciprocity requires firms to contribute to cooperative arrangements from which they benefit, but only when those arrangements are themselves fair, as Rawls argues is the case for individuals (Rawls, 1999, p. 96). And whether those arrangements are fair is the very matter at stake in international tax justice (Kern, 2020).

## **2.6 Chapter 2 Conclusion**

Existing normative theories of taxation can be roughly categorized according to the moral significance they afford some account of pretax holdings, and the extent to which they conceive of taxation as a distinct domain of justice. What I call the social architecture view focuses on the justice of distributive outcomes across a broad institutional framework. It emerges from dissatisfaction with the many variants of tax burden analysis, arguing that they are all subject to

the same criticism: they afford normative significance to an arbitrary baseline. That is, burden analysis takes pretax calculations as morally significant fixed points with which taxation interferes. Chapter one began to outline how the institutional framework powered by taxation in fact constructs those seemingly fixed points. Their general moral significance, if any, must depend on the justice of that framework, and so on justice in taxation already being secured. Indeed, those holdings would not exist at all without the state (and so taxation). This implication, that it is both morally and conceptually incoherent to claim general entitlement to pretax holdings, is the focus of the following chapter.

## Chapter 3: The Baseline Problem in Tax Fairness

### 3.1 Chapter 3 Introduction

This chapter considers the claim that there exists a baseline distribution of property rights with which taxation interferes. I call this commitment *the property constraint*. I then introduce Liam Murphy and Thomas Nagel's conceptual argument against the property constraint, the controversial first premise of which is that property rights are legal conventions. I briefly consider and reject a proposed alternative interpretation of their argument. I then observe that rejecting the property constraint is normally presented as a consequence of conventionalism. Therefore, I consider the meaning of 'convention' as it relates to property rights, leading me to distinguish between a convention and a *mere* convention.

With this distinction, I turn to the two main arguments against the first premise of the conceptual argument. The first argument is that property rights can exist outside and alongside the law, so it is false that all property rights are legal rights. I accept that the assertion of a property-law unity is either stipulative or empirically false. Therefore, I find this argument convincing: some property rights are not legal conventions. The second argument against the first premise of the conceptual argument is that all rights are protected by legal conventions, and so there is nothing exceptional about property. I note that classical liberal defenders of the property constraint conceive of property as a legal strategy for securing basic economic liberties. If the existence of legal conventions protecting property precludes a basic right to economic liberties, then so too must the existence of legal conventions protecting other basic liberties preclude a basic right to those liberties. But I do not think basic liberties are precluded by their realization in legal conventions. Therefore, I also find this argument convincing: the existence of legal property conventions does not preclude moral property rights.

Given that both arguments against the first premise of the conceptual argument succeed, the conceptual argument appears unsound. However, I submit that it may be rescued. It is saved if the first premise is reinterpreted so that it is trivially true: all *legal* property rights are legal conventions. This tautological interpretation is consistent with both critiques of the first premise, since it is not committed to the property-law unity and has no objectionable implications that would degrade other important rights. Moreover, this interpretation implies that the conceptual argument is independent of the kind of conventionalism many people find objectionable. I call this point *the independence thesis*. Given this tautological interpretation, the conceptual argument still concludes that legal property rights cannot provide a baseline distribution of entitlements against which the legitimacy of taxation can be assessed. Consequently, the conceptual argument achieves its purpose of demonstrating the incoherence of entitlement to the legal convention of some pretax distribution. However, this interpretation leaves open the possibility that *moral* entitlements, including moral property rights, can constrain taxation. But this is not surprising, since it is plausible that moral property rights do constrain the law, including the law of taxation.

### **3.2 Taxation and the Property Constraint**

As discussed above, the common sense of taxation assumes that individuals have at least provisional pretax entitlements to specific bundles of resources. Taxes reduce those resources by coercively taking them from individuals and transferring them to the state. On this view, taxation is a coercive interference with at least provisional entitlements and so the natural implication is that it demands compelling justification. This understanding of taxation is ubiquitous, often considered so obvious as to warrant implicit assumption. The very terms of tax analysis reflect

its ubiquity: there are *taxpayers* and *tax bills* and *tax burdens* and *redistributive policies* and so on. Often times it is also stated explicitly. Here are some examples:

All taxes infringe on rights, and to justify the imposition of taxes, one needs to demonstrate that the infringement is justified (Narveson, 2010, p. 128).

[T]axing requires justification . . . because it intrudes on the personal rights of those taxed (Vogel, 1988, p. 56).

Taxation is the power to coerce other individuals to surrender their property *without* their consent (Epstein, 1986, p. 49).

[T]axpayers or citizens . . . suffer the burden of taxation or are the potential subjects of government's powers of fiscal exaction (Brennan & Buchanan, 2000, p. 3).

[Taxed individuals are] deprived of their resources (de Jasay, 1997, p. 91).

[Taxation involves] exacting resources from [people], by force if necessary, in order to get them to do as they should (Ripstein, 2004, p. 5).

[S]ince my pretax income is something I have a right to, any infringement of that right must be justified (Gaus, 2010, p. 259).

Theft, taxation, and disembowelment are different forms of the same kind of violation of rights (Wheeler, 2000, p. 189).

Examples of this kind could be multiplied endlessly. Their authors are united in the conviction that taxation powers are exercised against a baseline distribution of property entitlements. Even those with more egalitarian sensibilities are normally inclined to accept the general framing of taxation as an interference with property rights:

Typically, any argument for either egalitarianism or redistribution of income through taxation takes the form of an argument for why property rights ought to be overridden for these purposes (equity or the public good) (Christman, 1991, p. 30).

This all suggests that thinking about taxation is heavily influenced by the assumption that a baseline distribution of property rights in some way constrains taxation. Variations of “some way” are myriad, but at its most general, call this assumption *the property constraint*.

In the last fifty years, however, a modest number of theorists from economics, law, and philosophy have expressed doubts about the property constraint. Their doubts arise from a similar insight which seems to have been developed at least four times independently.<sup>44</sup> Stated most generally, their skeptical insight coalesces around the assumed anteriority of property to taxation. As the introductory quotations indicate, the assumption of anteriority is deeply entrenched in the status quo. The first of the four skeptics to challenge the status quo was the economist Carl Shoup (1969) who asserted that the standard approach to assessing taxation “is conceptually invalid because it postulates, for implicit comparison, a state of affairs in which there are no taxes whatever. . .” (p. 577). Over two decades later, Charles Sampford and David Wood (1991) assailed the “priority of property,” lamenting that “taxation is always conceived of [as] taking from those who already have property, assets and income” (p. 181). John Christman (1994) then followed, criticizing the assumption that “application of the principles of economic justice are *imposed by* the state upon the individual citizen, tightening the circle that her *prima facie* property rights would have carved out” (p. 9). Finally, Liam Murphy and Thomas Nagel (2002) brought skepticism of the property constraint mainstream, arguing that “[i]t is illegitimate to appeal to a baseline of property rights in, say, ‘pretax income,’ for the purpose of evaluating tax policies, when all such figures are the product of a system of which taxes are an inextricable part” (p. 9). In short, these theorists all maintained that “[t]he logical order of priority between taxes and property rights is the reverse” of that implied by the property constraint (Murphy & Nagel, 2002, p. 33).

It is possible to conceive of the property constraint as posing an absolute prohibition on taxation. This is the often mentioned but less often defended view that ‘all taxation is theft.’ I

---

<sup>44</sup> Which is only to observe that none of the authors cite any of the others.

note this view only to set it aside, and to emphasize that for now I am interested in whether property rights pose a provisional, or *prima facie*, constraint—a justificatory hurdle of some significance, but not an impassable one. Approaching the question in this way seems justified because even devoted supporters of the property constraint, such as Geoffrey Brennan, describe ‘all taxation is theft’ as “a proposition that some radical libertarians claim to believe. Of course, almost nobody else does. In fact, it strikes most of the rest of us as ludicrous” (2005b, p. 130). At the same time, leading critics of the property constraint, such as Liam Murphy, describe that “radical libertarian proposition” as “an extreme view with little importance in contemporary politics” (2005, pp. 147-148). Since representatives of both sides of the debate describe an absolute constraint as extreme or ludicrous, and think its consideration of little importance, I trust it is reasonable to generally ignore it.

### **3.2.1 The Conceptual Argument**

The most promising place to begin is with an argument that aspires to decisively refute the property constraint. Recall that, understood most generally, the property constraint is simply the proposition that a baseline distribution of property rights in some way constrains taxation. If there is a persuasive conceptual or logical argument against the property constraint in its most general form, then this provides a conclusive reason to accept that a baseline distribution of property rights cannot pose even a provisional constraint on taxation. This may seem an unlikely prospect, but Murphy and Nagel (2002) claim to provide such an argument.

They characterize the target of their argument as the view that, absent redistributive government policy, “the pretax distribution of resources can be regarded as the distribution

produced by a free market” and is therefore morally significant (p. 32). But that assumption, they claim, “is deeply incoherent” (p. 32). This is their argument why:

There is no market without government and no government without taxes; and what type of market there is depends on laws and policy decisions that government must make. In the absence of a legal system supported by taxes, there couldn’t be money, banks, corporations, stock exchanges, patents, or a modern market economy—none of the institutions that make possible the existence of almost all contemporary forms of income and wealth (p. 32).

The argument is simple, perhaps deceptively so. It is simply that property rights are a creation (in part) of taxation and so cannot provide an independent baseline of entitlements against which to assess the justice of the tax system.

All [people] can be entitled to is what they would be left with after taxes under a legitimate system, supported by legitimate taxation—and this shows that we cannot evaluate the legitimacy of taxes by reference to pretax income. Instead, we have to evaluate the legitimacy of after-tax income by reference to the legitimacy of the political and economic system that generates it, including the taxes which are an essential part of that system. The logical order of priority between taxes and property rights is the reverse of that assumed by libertarianism (p. 33).

This grounds their claim that any notion of entitlement to all of pretax income is incoherent.

Though they never express the argument formally, they repeatedly emphasize that it is intended to establish a logical or conceptual point:

It is therefore logically impossible that people should have any kind of entitlement to all their pretax income (2002, p. 32).

[A denial of collective ownership] doesn’t justify starting with pretax income—over which individuals *couldn’t, as a matter of logic*, be given full private control—as the baseline from which departures must be justified (2002, p. 176).

For whatever the truth about the morality of property, we couldn’t possibly have a right to the whole of our pretax income. The point here, unusually in practical philosophy, is one of logic (Murphy, 2005, p. 124).

This logical point is entirely general: Whatever we believe we are entitled to as a matter of distributive justice or property theory, or anything else, we cannot be entitled to the whole of our pretax income (Murphy, 2005, p. 124).

In order to better assess the thesis being expressed in these and other statements, this is an attempted reconstruction of what I take to be *the conceptual argument*:

1. Property rights are legal conventions.
  2. Legal conventions are created (in part) by taxation.
  3. Therefore, property rights are created (in part) by taxation.
  4. It is logically impossible for the creation (in part) of an institution X to provide an entitlement baseline independent of X for assessing the legitimacy of X.
- 
5. Therefore, it is logically impossible for property rights that are the creation (in part) of the institution of taxation to provide an entitlement baseline independent of taxation for assessing the legitimacy of taxation.

Something like this argument is latent in all the skeptical assessments of the property constraint.

Another example comes from Sampford and Wood (1991) who see defenders of the property constraint as right-wing versions of the “vulgar Marxists who insist on a rigid determination of the Superstructure by the Base” (p. 182). The mistake is fundamentally the same: both assume an economic structure prior to law when that structure “cannot exist without the law” (p. 182). The “vulgar capitalists” simply “assume the conceptual priority of pretax holdings and make no attempt to justify their moral priority” (p. 183). Critics of the conceptual argument against the property constraint have, of course, provided a number of responses. These generally aim to reject the first premise and are taken up below. First, however, one critic suggests a different interpretation of the statements I referred to in order to construct the conceptual argument. Geoffrey Brennan calls his interpretation of Murphy and Nagel’s logical point the “incoherence thesis” (2005a; 2005b; 2018).

### **3.2.2 The Incoherence Thesis**

Brennan suggests that the conceptual argument is quite different from my proposed reconstruction. Instead, he thinks Murphy and Nagel intend to make a logical critique of the very

notion of gross-of-tax magnitudes. In order to assess this possibility, he calls it the “incoherence thesis” and describes it as the thesis that gross-of-tax magnitudes are conceptually incoherent. But he doubts the incoherence thesis is plausible; in fact, he thinks support for it reveals a “confusion as to the special status of ‘pretax’ magnitudes” (2005a, p. 212). And he suspects such confusion is evident in Murphy and Nagel’s failure to realize that pretax income *just is* after-tax income plus taxes paid. That is, the notion of pretax income as they describe it is equivalent to the economists’ gross-of-tax magnitude, which is simply the sum of the plainly coherent ideas of after-tax income and taxes paid. Indeed, the formulations are two sides of the same coin. It is true that these figures exist in mutual interdependence with public spending, he continues, but this “is not sufficient to show that a tax regime that makes reference to these purely conceptual magnitudes is, ‘*as a matter of logic*’ inconceivable” (2005a, p. 212). As Brennan notes, “[a]ppeal to gross-of-tax magnitudes is just an accounting detail: there is nothing ‘illogical’ or ‘incoherent’ about that aspect of the argument” (2005a, p. 212). If Murphy and Nagel think otherwise, he finds them “guilty of an error – that of mistaking an arithmetic convenience for a conceptual claim” (2005a, p. 212).

Brennan is right that accepting the incoherence thesis, as he describes it, would be a mistake. Pretax income is just one species of gross-of-tax magnitude, and so both terms can express the same coherent idea. Therefore, if Murphy and Nagel accept that “after-tax plus taxes paid” is a coherent calculation, they cannot reject the coherence of “gross-of-tax less taxes paid.” But there are two important observations about Brennan’s proposed interpretation. The first is that Murphy and Nagel do not appear to argue that a tax regime that “*makes reference to* [emphasis added]” pretax income is inconceivable (Brennan, 2005a, p. 212). On the contrary, they readily acknowledge that pretax magnitudes provide pragmatic “book-keeping figure[s]” for

a system of taxation (2002, p. 36, 75). The second (and related) important observation is that Brennan's construal of the incoherence thesis seems to miss the essential point that the conceptual argument concludes that any *entitlement* to all of pretax income is incoherent. His oversight is especially clear in the following statement of the pretax / gross-of-tax symmetry:

[I]f tax norms are such as to specify that tax liabilities have to be equal for individuals with identical pretax incomes, this is exactly equivalent to the requirement that tax liabilities be identical for those with identical net-of-tax incomes (or whatever the tax base is) (2005a, p. 212)

This is of course unobjectionable, since it only points out that the arithmetic can be expressed in a different way. But it overlooks that the "if" is precisely what is at stake in the argument. Critics of the property constraint doubt that pretax income (or, if Brennan prefers, gross-of-tax magnitudes) can provide a baseline *entitlement* against which "equal" liabilities might be assessed. The conceptual argument concludes that property rights in pretax income cannot coherently provide such a baseline. Brennan's observation is simply not responsive to that argument. In summary, Brennan rightly rejects the incoherence thesis (as he describes it), but wrongly attributes it to Murphy and Nagel, and fails to respond to the conceptual argument against the property constraint.

### **3.2.3 Property Rights and Legal Conventionalism**

Most of the debate about the conceptual argument centers on the first premise, that all property rights are legal rights. This is understandable since that proposition is both ambiguous and controversial. Moreover, critics of the property constraint generally treat the assertion that property is a legal convention as central to their view. For example, Sampford and Wood (1991) want to distance themselves from "the utterly discredited view that property is pre-social" (p. 201). As later discussed, legal conventionalists routinely defend their position with examples that

illustrate how property in the modern world seems inextricably bound up with law. And of course legal conventions are not natural, but are a paradigm case of social practice:

[A]ny current distribution of property . . . is a complex outcome of the combination of legal rules regulating contracts, employment, inheritance, familial division of property, taxation, torts, expropriation or confiscation with or without compensation etc., and of human actions undertaken in the context of this pre-existing legal regime. There is nothing ‘natural’ about these rules in that they represent the value judgments and policy decisions of law-makers; they enforce public choices which are inherently controversial (for any particular regime of the tax law one can propose an alternative one, which will result in a different distribution of wealth, etc.). These rules are just as coercive as the rules which govern the use of force against violators of the status quo: no more and no less (Sadurski, 1989, p. 465).

Murphy and Nagel (2002) also place legal conventionalism about property at the heart of their analysis: “[t]he conviction that determines our approach to all more specific questions is that there are no property rights antecedent to the tax structure” (p. 74). They call legal conventionalism the “dominant theme” of their discussion (p. 8). They emphatically deny that “property rights are morally fundamental” (p. 175), claiming instead that it is “perfectly obvious” (p. 8) and “common sense” (p. 74) that “private property is a legal convention” (p. 8). As a result, they enthusiastically embrace a “purely conventional conception of property” (p. 175) and assert that “individuals don’t own anything except through laws” (p. 176).

The idea that property rights are “mere conventions” (Murphy and Nagel, 2002, p. 73) is part of a larger tradition that is often traced to the thought of Jeremy Bentham (1864) or, as referenced above, David Hume (1739).<sup>45</sup> The general idea is that property rights are simply a tool which may (or may not) be instrumental in achieving other valuable ends. Private property in particular is a legal creation and not a structure of rights independent of law to which the law

---

<sup>45</sup> Instrumental accounts of property are, of course, much older. For example, I noted above that Aristotle articulates a number of instrumental arguments for private property in Book II of his *Politics*.

should conform. The following provides some representative statements of the kind of instrumentalist view I have in mind:

[P]roperty is a legally constructed social relation (Holmes & Sunstein, 2000, p. 59).

[P]roperty is legal possession (Kekes, 2010, p. 3).

[Property rights are] an instrument of society (Demsetz, 1967, p. 347)

[P]rivate property is basically just a tool (Freyfogle, 2007, p. xxi)

We *create* private property institutions in service of a variety of fundamental, though perhaps lower-level, goods (creator-incentive, for example) for both individuals and society, and we design such institutions using the tools (a market, for example) which will best serve these goals (Lametti, 1998, p. 715).

Indubitably, property is a political creation. After all, we Americans have watched the creation of privatised property rights for decades: intellectual property rights over the Internet, over gene sequences, and over delivery of health services to the poor and the elderly (Frajman & Oppenheimer, 2003, p. 385).

Conventionalists of this sort are pleased to accept the first premise of the conceptual argument.

As a result, they are understandably mystified by the suggestion that some extralegal distribution of property rights can act as an independent *moral* constraint on taxation; after all, property rights and taxation are both just *legal* practices on this view.

Insisting on the conventionality of property rights suggests a conditional relationship. The thought is that if, or even *because*, property is conventional, there can be no property rights to pretax income. If this conditional relationship holds, then vindicating (or defeating) conventionalism is crucial for assessing the conceptual argument. And indeed, this is generally how the argument has been understood. For example, Miranda Stewart claims it is “axiomatic to [Murphy and Nagel’s] discussion of tax justice [. . .] that private property is a legal convention” (2005, p. 133). Robert Young assesses the argument as depending on the thesis that there is “no natural right to property” (2005, p. 143). The rhetorically implied upshot of this conditional

relationship is that if property is *not* merely conventional (extralegal ownership *not*, after all, a myth), then entitlement to the pretax baseline is vindicated.<sup>46</sup> In the worst case, therefore, “to show that property is not mythical would be to demonstrate that taxation *is* indeed theft” (Brennan, 2005a, p. 208). In general, the conditional relationship (if true) does imply that defeating conventionalism is necessary for even the *possibility* of redeeming property rights in some pretax distribution.

So far we have seen that one tactic for rejecting the property constraint is to endorse the conceptual argument. In turn, the conceptual argument depends on conventionalism (or so it is thought). Therefore, defenders of the property constraint—and so critics of the conceptual argument—have taken two general approaches in their attempts to undermine the conventionalist premise that all property is law. The first is to show how property exists both outside and alongside the law and thus cannot be identical with it. The second is to show how conventionalism fails to illuminate anything exceptional about property rights. As I will explain, I think both of these approaches are generally correct. However, I do not think they defeat the conceptual argument, properly understood. Before evaluating the two approaches, we need to clarify what conventionalism means in this context.

### 3.2.4 Conventions

What is a convention? The first thing to note is that the concept is ambiguous:

The terms “convention” and “conventional” are flagrantly and intricately ambiguous. On the one hand, the conventional is the ordinary, the usual, the traditional, the orthodox as against the novel, the deviant, the unexpected, the heterodox. On the

---

<sup>46</sup> I say it is “rhetorically implied” (rather than logically implied) because, as I later discuss, drawing this conclusion is a logical mistake (denying the antecedent), and has important consequences. If conventionalism implies no property rights in pretax income, then conventionalism must be defeated for even the possibility of establishing property rights in pretax income. However, defeating conventionalism does not, in itself, establish those property rights. To conclude otherwise, as Brennan’s framing suggests, is a case of denying the antecedent.

other hand, the conventional is the artificial, the invented, the optional, as against the natural, the fundamental, the mandatory (Goodman, 1990, p. 93).

In his recent survey of conventionalism, Michael Rescorla finds that “one can say little of substance about it as a general matter” (2019, 1.2). An initial thought is that conventions imply some kind of agreement, but then paradigm cases of linguistic and legal conventions are manifestly not cases of agreement. Rescorla finds that if there is one “distinctive thesis shared by most conventionalist theories” it is that “there exist *alternative* conventions that are in some sense equally good” (2019, 1.2). The existence of equally good alternatives seems to accord well with conventions of language where “‘book’ could just as easily refer to what ‘plant’ refers to, and vice versa” (Smith, 2003, p. 1106). In the property context, Murphy and Nagel also make an analogy to linguistic conventions. They point out that it would be senseless to find a pig rightly named “on account of its eating habits and tendency to wallow in the mud” (2002, p. 74). The meaning of “pig” is a convention that cannot be used to justify itself.

The availability of equally good alternatives suggests that conventions are, in some important sense, arbitrary. Indeed, in their discussion of the pig, Murphy and Nagel also emphasize that linguistic conventions are “highly arbitrary” (2002, pp. 74-75): a pig might just as well have been called something else; its name is only a convention. In fact, in his influential treatment of the concept, David Lewis finds that *all* conventions are arbitrary. He even claims that the phrase “arbitrary convention” is redundant:

If R is our actual convention, R must have the alternative R', and R' must be such that it could have been our convention instead of R, if only people had started off conforming to R' and expecting each other to. This is why it is redundant to speak of an arbitrary convention. Any convention is arbitrary because there is an alternative regularity that could have been our convention instead (1969, p. 70).

Lewis' point about arbitrariness seems to illuminate an important feature of some conventions.

For example, people should coordinate how they drive on the road. Yet whether they drive on the

left or on the right is (I assume) arbitrary, a matter of convention. But then other cases fit less well with this analysis. A speed limit seems like a convention, but it is not arbitrary. In fact, even the case of driving on the left or the right is only arbitrary because the choice set is constrained to include only equally good alternatives. We could include a third alternative, in which the bidirectional traffic rules switch every three hours: everyone drives on the right for three hours, then on the left for three hours, and so on. This would be a convention, but one with powerful normative considerations (I assume) against adopting it. Therefore, either some apparent conventions are not conventions in Lewis' sense after all, or Lewis is wrong that all conventions are arbitrary (see, e.g., Jamieson, 1975).

Another potential feature of conventional claims is that conventions are all that makes them true. This is especially relevant here because, as Rescorla observes, "Murphy and Nagel seem to argue . . . that claims about property rights are *made true* by convention" (2019, fn. 2). This also seems to track Jan Narveson's analysis of legal ownership: the "case for A's being taken to be the legitimate owner of  $x$  . . . may be political and legal in the sense that the question is whether A's possession of  $x$  is recognized and upheld in law or in the political structure of the community in question" (2010, p. 107). On this view, the proposition 'A has legal property rights in  $x$ ' is made true (or false) by legal convention. As with arbitrariness, this seems to illuminate an important feature of some conventions. If it is true that people should drive on the right, then the convention to do so is what *makes it* true. If such a convention exists then, precisely *because* it exists, people should drive on the right. By implication, the normative force of the 'should' is entirely dependent on the convention, and so it recedes in the absence of the convention (Brennan, 2005a). In case a new convention emerges to drive on the left, people no longer have any reason to drive on the right. Again, however, other cases fit less well. Even in the absence of

a convention to drive a reasonable speed, there remains good reason to do so. This seems especially clear when considering legal conventions prohibiting, for example, murder (Brennan, 2005a).

This might suggest a distinction is necessary. Perhaps there are conventions (e.g., speed limits) which are not arbitrary, and then there are conventions (e.g., driving on the right), which are arbitrary. Or perhaps there are conventional claims that are made true *only* by convention (e.g., a legal requirement to drive on the right) and conventional claims with normative force independent of convention (e.g., a law prohibiting murder). The natural distinction might therefore be made between on the one hand conventional claims that are nonarbitrary and have independent normative force, and on the other hand *merely* or *purely* conventional claims that tend towards arbitrariness and, if they are true, are true only because of a convention.

Legal property rights are conventions, but they are not *mere* conventions in my proposed sense: they are not arbitrary and they do not float free of independent normative force. No one believes that any allocation of property rights is equally as good as any other. A choice among property conventions is not like a choice between driving on the left or driving on the right; the property choice must be assessed on normative grounds. Indeed, Murphy and Nagel (2002) invoke the longstanding conventions of slavery (p. 8) and discrimination against women (p. 9) precisely to show that conventions are often not arbitrary but are subject to moral evaluation. Relatedly, as I will later review, there is good reason to think that powerful normative considerations ground property claims quite independent of any legal convention. For one example, Jewish art was legally stolen by the Nazis (O'Donnell, 2011). Yet the conceptual possibility of legal theft poses a serious challenge for the *mere* conventionalist.

Despite the conspicuous fact that legal property rights are not *mere* conventions, ambiguity about the meaning of “convention” has proved fertile ground for confusion in assessments of the conceptual argument. How can this be explained? Murphy and Nagel are at least partly to blame. They have invited a misunderstanding of conventionalism, since they compare their view to “highly arbitrary” (p. 74) linguistic conventions, and they are emphatic that property rights are *purely* (p. 175) and *merely* (p. 73) conventional. Yet they never distinguish between a convention and a *mere* or *pure* convention. The most plausible grounds for that distinction (arbitrariness, independent normative force) imply on their own account that property rights are *not* mere or pure conventions after all.

The resulting confusion manifests in the common but mistaken idea that conventionalism implies a kind of open season on property rights. In short, the thought is that if property is a *mere* convention, and *mere* conventions are arbitrary, then any distribution of property is as good as any other. This is evident in Brennan’s suggestion that conventionalism implies that whatever a person “can legitimately lay claim to is, well, whatever the law says she can lay claim to” (Brennan, 2005, p. 130). He is not alone in drawing this conclusion. Gerald Gaus characterizes Murphy and Nagel’s argument as holding that “since property rights are purely conventional, the state may determine their shape *as it sees fit*, [emphasis added] and this includes determining the level of taxation” (2010, p. 263). Though he agrees that “property rights are conventional and depend on the state” (2010, p. 260), he points out that this does not imply that there is “no Archimedean point outside the state’s determinations of your property rights (or any other

rights?)<sup>47</sup> from which to criticize the state's activity as taking what is yours; for its decisions determine what is yours" (2010, p. 260). Likewise, Jeffrey Schoenblum excoriates what he sees as the conventionalist view that "since the taxpayer is not 'entitled' to *anything*, there are no constraints on how much the government can take" (2003, p. 560). In the end, he finds conventionalism equivalent to the dismal conclusion that "[t]he 'government' is the true owner of material wealth" (Schoenblum, 2003, p. 557).

These critics of conventionalism miss the mark. As noted, describing property rights as merely and purely conventional invites the distorting inference that conventionalism has "denuded the individual of any moral claim to private property" (Schoenblum, 2003, p. 569). However, Murphy and Nagel explicitly reject the claim that "the entire social product really belongs to the government, and that all after-tax income should be seen as a kind of dole that each of us receives from the government, if it chooses to look on us with favor" (2002, p. 176). Indeed, Murphy later emphasizes that "nothing could be further from the spirit of the book than the idea that whatever the government does, through law, in the realm of property and economics, is by definition just" (2005, p. 149). Even for Murphy and Nagel, then, property rights are decidedly *not* mere conventions in either of the two plausible senses I described above.

---

<sup>47</sup> The mistaken suggestion here is that conventionalism implies that dependence on the state undermines the status of rights. It also has some traction. For instance, in Schoenblum's words: "the authors present the following flawed logic: (a) the collectivity creates conditions, especially through law, enabling the production and preservation of wealth; (b) the government *is* the collectivity; and (c) the government, therefore, is entitled to control *all* of the wealth that exists and is created (2003, p. 562)." He continues, "[f]urther, the implications of the authors' reasoning extend far beyond private property rights. If, in fact, the government can claim 100% of all wealth because it makes private property possible, then it can also circumscribe or deny all civil rights that it also makes possible through the enforcement of law (2003, p. 563).

Timothy Fowler (2015, p. 349) interprets John Tomasi (2012) as leveling a similar charge against the conventionalists: "Murphy and Nagel are committed to the absurd suggestion that if tax laws cannot violate property rights then no laws can ever violate citizens' basic rights."

We can set aside any exegetical concerns<sup>48</sup> and proceed with the following understanding of conventionalism and property. Legal property rights are conventions. But legal property rights are not arbitrary, and (as with criminal laws) their status as legal conventions does not imply that they have no normative force independent of law. The consistency of these ideas can be illustrated with a simple example of a classical conventionalist. Jeremy Bentham rejects the natural rights tradition and conceives of legal property as conventional. Yet he does not think any arrangement of property rights is as good as any other. The correct arrangement also has normative force independent of law since it is dictated by the maximization of utility as a master value. It is to Bentham's expression of the first premise of the conceptual argument that I now turn.

### **3.2.5 The Property-Law Unity**

The first premise of the conceptual argument is that property rights are legal rights. As the examples above show, "it is common to say that property is an institution created by the state" (Merrill, 2012, p. 2076). But perhaps nowhere is this more forcefully expressed than in Bentham's famous statement that "[p]roperty and law are born and must die together. Before the laws, there was no property: take away the laws, all property ceases" (1864, I.VIII). Critics of this conventionalist conception of property have sought to demonstrate that, contra Bentham, property exists both outside and alongside legal institutions. If property rights are not only legal rights, then the first premise of the conceptual argument is false, and the argument fails. Or so it is thought.

---

<sup>48</sup> All I mean to set aside here is Murphy and Nagel's particular meaning of "mere" and "pure" as qualifications that attach to certain conventions. I have proposed what I take to be meaningful interpretations of those qualifiers, but my proposal does not track their usage. In fact, I doubt their usage of the terms means anything of substance. I think the qualifiers are deployed as hyperbolic adjectives to emphasize their rejection of the natural rights tradition.

We should concede that asserting all property rights are legal rights is either stipulative or implausible. In fact, despite his frequent citation on the necessary relation between property and law, I suspect even Bentham realized it was not really true. He defines property as a settled collection of expectations: a relationship to an object that is recognized by others. But of course law is not the only way to settle expectations. He tries to evade this problem by describing all apparently non-legal conventions as legal principles. He illustrates the point with a “savage” who in a “primitive state. . .has hidden his prey” and can only rely on his own strength to protect his bounty from rivals (1864, I.VIII). Bentham accepts that a truce could arise in such circumstances, but then holds that (conveniently) even “the slightest agreement” amounts to a legal principle (1864, I.VIII). Even then, he hastens to add, such a law is only a “feeble and momentary expectation” in contrast with the “strong and permanent expectation” found in a society (1864, I.VIII). Thus, his apparently conceptual claim is actually empirical—that only law can provide the clarity and stability of expectations necessary for private property.

The problem for Bentham-type arguments, then, is simply the overwhelming anthropological evidence that the key empirical premise is not true. Anthropologists have observed that “property rights are a critical dimension of the evolution and materialization of social institutions and political relationships” (Earle, 2000, p. 53). But this is qualified by acknowledging that “*codified laws of property can be considered as but one mechanism* [emphasis added] through which objects and land relate behaviorally to people's use, allocation, and transfer. Property is integral to all concepts of social institutions—how people are related to resources and to each other” (Earle, 2000, p. 40). Law is but one way to settle expectations, and to mediate relations among people and resources. Moreover, there is satisfactory evidence that “notions of property pre-date the establishment of modern systems of law and enforcement”

(Brennan, 2005a, p. 211; see, e.g., Hoebel, 1954) and also of “[s]ocieties [that] have existed with property and without the institution of taxation” (Frajman & Oppenheimer, 2003, p. 387; see, e.g., Bailey, 1992). Even Murphy accepts that “property in societies that lack formal institutional structures backed by centralized coercive force” is a “conventional normative practice” (2005, p. 48). A conventional normative practice, however, is not necessarily law. Therefore, it is false that all property ceases in the absence of law, unless property is stipulated to mean legal property rules. Bentham’s assertion “take away the laws, all property ceases” is therefore empirically false or a trivial tautology.

There is also the problem that property relations are not exhausted by the legal system. In fact, property law itself is “infused with moral value” (Merrill & Smith, 2007, p. 1870) and practices at its core (exclusion, alienation, use) are governed by social rules. Thus critics are understandably skeptical of “the assumption that property is purely a creature of law—in [Murphy and Nagel’s] case, the tax system—and has no other force, presumptive or otherwise” (Merrill & Smith, 2007, p. 1870). A social practice of property exists both outside and alongside the law. And this practice is distinctively *of property* which is not well-explained by its promotion of other values. To see property operate alongside the law, there are numerous case studies of communities with property practices and norms that ignore, or even conflict with, the law. For example, Robert Ellickson’s famous study of rural landowners in Shasta County illustrates how informal norms of trespass (not the law of trespass) guide interactions:

The field evidence I gathered suggests that a change in animal trespass law indeed fails to affect resource allocation, *not because transaction costs are low, but because transaction costs are high*. Legal rules are costly to learn and enforce. Trespass incidents are minor irritations between parties who typically have complex continuing relationships that enable them readily to enforce informal norms. The Shasta County evidence indicates that under these conditions, potential disputants ignore the formal law (1986, p. 628; see also, Ellickson, 1994; Acheson, 1988).

A morality of property also influences the law from within (Merrill & Smith, 2007). For instance, trespass rules are often vigorously upheld seemingly for their own sake. Merrill and Smith illustrate this with the representative case of *Jacque v. Steenberg Homes*.

The case involved a dispute between an elderly Wisconsin couple, Harvey and Lois Jacque, who lived on 170 acres of rural land near Schleswig, Wisconsin; and a mobile home sales company, Steenberg Homes, which wanted to deliver a mobile home to a neighbor of the Jacques (2007, p. 1871).

Steenberg offered to pay the Jacques for temporary access, since the neighbour's access was blocked. The Jacques refused, citing unfounded fears of exposure to adverse possession.

Steenberg cut across the land and delivered the home anyway. The court's response was notable for imposing damages on Steenberg "one hundred times the largest criminal fine that could be imposed for this behaviour" and making "no attempt to consider whether the trespass would have been cost-justified" (p. 1872). No reasonable person would deny access, and any comparative utility considerations were fully on the side of Steenberg. The court nevertheless ruled in "deontological terms" and "did not regard the justification for punitive damages to be reducible to nonproperty values" (p. 1873).

There is also the fact, evident in everyday experience, that people do not act as though the contours of property are fixed by the law. People routinely recognize and respect one another's claims to use, exclude, and alienate quite independently of law. Legal sanctions are often weak and poorly enforced, yet people tend not to trespass or steal. One might think this reflects a general moral commitment to follow the law. But such a commitment fails to explain why, when common morality diverges from law, there is widespread disregard of legal property rules. This is evident, for example, with downloading copyrighted materials (Merrill and Smith, 2007) or the return of legally stolen art (O'Donnell, 2011). No doubt law matters in structuring all these practices, but there is little reason to believe that it is *all* that matters. As Armen Alchian puts it,

property is “supported by the force of etiquette, social custom, ostracism, and formal legally enacted laws” (1977, p. 129). There are distinctive property norms and practices external to the law and interwoven with it.

This appears to be a significant victory in favour of the property constraint. The most powerful argument against it is the conceptual argument, and the first premise of that argument is either trivial or implausible. But how serious is this dilemma? Defenders cannot countenance an implausible premise since that would be fatal to the success of the argument. A trivial premise, on the other hand, need not undermine the conceptual argument. A trivially true or tautological premise does not prevent the argument from going through. In fact, as I will argue, the premise *should* be understood as trivially true: all *legal* property rights are *legal* conventions. Far from defeating the argument, this interpretation reveals that it is more powerful, and more general, than is commonly appreciated. But first I consider the alternative approach to rejecting the conceptual argument.

### **3.2.6 Economic Exceptionalism**

The first approach to rejecting the conceptual argument relies on undermining the truth of its first premise (that all property rights are legal rights) by defending property as a moral practice both outside and alongside law. That approach rightly demonstrates that there are property norms and practices that do not depend on law. I now turn to the second approach, which involves showing that conventionalist arguments purporting to prove that all property rights are legal rights show nothing exceptional about property. Once this is explained, I consider the implications of the successes and/or failures of the two approaches.

Attempts to establish the conventionality of property tend to press that some property rights seem *necessarily* legal—they cannot be understood any other way. They point out that any fantasy of a Lockean man plucking an apple from a tree is far removed from property in a modern economy. Instead of “unoccupied” land and apples, property rights address complex resources including those considered in chapter one: electromagnetic spectra, data, trademarks, fiat currency, patents, use rights (to fish, to mine, etc.), complex financial instruments, licenses, corporate shares, options, reputations, artistic expression, labour, debt, and so on. A world of land and apples is tangible and concrete, but in our world valuable resources are often intangible and abstract. Moreover, many of these resources seem legal by their very essence. Consider an incorporated holding company and its wholly owned subsidiary. Both owner and owned are legal entities, a relationship that is fundamentally constructed by law. The conventionalist then points out that law is a paradigm case of a convention, so clearly corporate property rights, for example, are necessarily legal conventions (see, e.g., Sampford & Wood, 1991).

There is one sense in which this claim must be right, but it is less clear what it is supposed to demonstrate. Only the radical libertarian could doubt that fundamental moral rights are normally realized in law. Certainly, defenders of the property constraint do not hesitate to depart from libertarianism on that point. Friedrich Hayek’s assertion is unremarkable in that regard: “recognition of the right of private property does not determine what exactly should be the content of this right” (1960/2013, 338). This is later echoed by Milton Friedman, who agrees that “what constitutes property and what rights the ownership of property confers are complex social creations rather than self-evident propositions” (1962, p. 26). Modern defenders of the property constraint, such as John Tomasi, concur: “[a]ll basic rights and liberties are socially constructed in important ways” (2012, p. 71).

That there is an important element of social construction in the legal protection of even the most basic rights and liberties seems like common sense. Consider any of the basic liberties such as the right to a fair trial, free association, free expression, equal treatment under the law, and so on. Can the abstract moral right to a fair trial specify the precise contents of the legal rights held by a criminal defendant? Is the defendant entitled to a jury, or only a judge, or a choice between the two? Do entitlements vary along with the severity of the charge? Should peremptory challenges be permitted in jury selection (*R. v. Chouhan*)? Or consider freedom of expression, where “contextualization and responsiveness to the requirements of other rights are assumed as a matter of course” (Breakey, 2014, p. 573). Then there is equality before the law, the precise contours of which seem especially dependent on the existence and character of the law. All the standard basic liberties are protected through a regime of legal conventions, and the specific content of those protections is inseparable from their particular social context.

Apparently the mere observation that there are legal conventions of property says nothing about the moral status of property rights. At minimum, as Tomasi points out, the argument to the contrary “is generalizable so as to render it either trivial or circular” (2012, p. 70). Then why do conventionalists routinely point to legal conventions of property as *evidence* that property rights are *merely* legal rights? Hugh Breakey (2014) suggests both an answer and a framework of analysis. His answer is that a number of theorists, including Murphy and Nagel, “conflate the core concept of a right with a situated or specified right from which one can read off people’s actual legal entitlements and duties” (p. 573). In making his diagnosis he analyzes the concept of a right, noting that the term ‘right’ is used to refer to five distinct notions.

Breakey’s central idea is that rights can be understood at varying levels of specificity, “from concept to conclusive” (2014, p. 578). His five notions begin at the most abstract with the

idealized “core right” which is a decontextualized account of the freedom the right aims to secure. Next is a “coherent right” which must be constrained such that it could be held by all, and it must not “foreclose the institutional requirements of its protection” (2014, p. 580). A “prima facie right” has additional content, as it limits direct harms and introduces Rawlsian-type “fundamental cases” in its “central range of application” (2014, p. 581). A situated right takes account of all other rights (not only rights of the same kind, which are accounted for in the coherent right). Finally, the specified right accounts for when, in certain conditions, the right may be superseded. He summarizes the relations between his five proposed notions:

The core concept of the right is *shaped* by the requirements of logical possibility to create the coherent right. The coherent right is *contoured* by paradigm cases and harmful cases to create the prima facie right. The prima facie right is *balanced* by the existence of all other relevant rights to create the situated right. The situated right, in special circumstances, may be *infringed* (but not violated) pursuant to the specified right. Any further intrusions on the specified right are *violations* of the right and are politically wrongful (2014, p. 583).

There is plenty of room for disagreement about the details, but the implication of Breakey’s central thesis is compelling: the core concept of a right cannot directly entail its situated form in a given set of institutions. According to Breakey, this explains how conventionalists go wrong.

The argument is roughly as follows. Some legal rights are intended to realize basic liberties, and those liberties have independent normative force. Other legal rights have no normative force independent of their status as law. But *both* kinds of legal rights entail specific legal conventions. So the “common-sense” observation that property rights are realized in legal conventions does not say anything one way or the other about whether they have independent normative force. Conflating the core right with the situated or specified right has, according to Breakey, caused theorists to miss this point, and therefore to erroneously conclude that the undeniable legal conventions of property demonstrate that property rights are *nothing but* legal

rights. This conclusion does not follow, and so Breakey concludes that property rights are not undermined by observations of legal conventions of property.

Once the conflation between core and situated rights is avoided, everyone (except the radical libertarian) looks like a conventionalist. Defenders of the property constraint certainly resist *mere* conventionalism, but then so do its critics. This explains why Murphy and Nagel (2002) can insist that conventionalism is “perfectly obvious” (p. 8) and “common sense” (p. 74) and yet, at the same time, Tomasi does not disagree: “[w]ith respect to assertions about the legitimacy of restrictions on economic liberty, the claim that such liberties are legal conventions is trivial” (2012, p. 71). Conventionalists point to legal practices (for instance, necessarily legal property) to defend the claim that all property rights are legal rights. Yet even their critics agree that there are necessarily legal conventions of property. The real source of disagreement lies elsewhere.

The views of Breakey and Tomasi illuminate the real source of disagreement: the moral status of economic liberty. Defenders of the property constraint are not radical libertarians who assume the legal system must perfectly track some pre-existing structure of moral property rights. That assumption depends on the “absurd” belief that “logical deduction can take us from the general principle of self-ownership to detailed legal arrangements” (Fried, 2004, p. 71). Instead, defenders of the property constraint think economic freedom is a powerful normative consideration, though how it is embodied and codified in law (as in social practice) is obviously a matter of convention. In fact, this disagreement is at the very core of Samuel Freeman’s (2011)

distinction between classical<sup>49</sup> and high<sup>50</sup> liberalism (see also Schumpeter, 1954). A number of considerations motivate Freeman to draw this distinction, but disagreement over the moral status of economic liberties is fundamental:

Where liberals primarily disagree is on the nature and status of economic rights and liberties, including the extent of freedom of contract and rights to private property in land, raw materials, and other productive resources (2011, p. 20).

For classical liberals, protecting private property and extensive freedom of contract are the main ends of civil society. As Gerald Gaus describes it:

On the classical view, the main tasks of the justified state are to protect individual freedom and secure a regime of extensive private property rights—that is, an economic system in which owners have an extensive bundle of rights over their property and in which the range of resources and assets subject to private ownership is extensive (2010, p. 235).

But with the high liberalism of John Rawls, these commitments suffer a certain fall from grace.

In his *Tanner Lectures*, Rawls includes *personal* property in his account of basic liberties:

Among the basic liberties of the person is the right to hold and to have the exclusive use of personal property. The role of this liberty is to allow a sufficient material basis for a sense of personal independence and self-respect, both of which are essential for the development and exercise of the moral powers (1981, p. 12).

But he echoes his earlier exclusion of extensive economic liberties of contract, acquisition, bequest, and control over productive resources:

Of course, liberties not on the list, for example the right to own certain kinds of property (e.g., means of production) and freedom of contract as understood by the doctrine of laissez-faire are not basic; and so they are not protected by the priority of the first principle (Rawls, 1971, p. 54).

---

<sup>49</sup> Representative theorists from Freeman (2011, p. 20): [1700s] Adam Smith [1800s] David Ricardo, Thomas Malthus, J.B. Say [1900s] Friedrich Hayek, Milton Friedman, James Buchanan, Gary Becker, Ludwig von Mises, David Gauthier, Richard Posner, Richard Epstein

<sup>50</sup> Representative theorists from Freeman (2011, p. 27): [1800s] John Stuart Mill. T.H. Green, John Dewey [1900s] John Rawls

His reasoning for this demotion vis-à-vis the classical tradition, a demotion which Tomasi terms “economic exceptionalism” (2012, p. 42), is that, in Rawls’ view, extensive economic liberties are not “necessary for the development and exercise of the moral powers” (1981, p. 12).

Predictably, this is far from convincing for those with classical liberal sympathies, seeming more like an assertion than an argument (e.g., Tomasi, 2012).

I do not mean to take a position on the high liberal exclusion of the classical economic liberties from the set of basic liberties. I do, however, think that exclusion is the real point of difference between conventionalists and their critics. This is why I think criticisms of the property constraint that point to the “perfectly obvious” fact that property rights are legal conventions, and then conclude that all property rights are *only* legal rights, are unlikely to succeed. After all, defenders of the property constraint simply believe those legal conventions can and should protect some prior basic liberty. As I noted before, both sides evidently agree that property rights are legal conventions. The question of substance is whether those conventions protect a fundamental freedom or ‘core’ right, as is the case with the legal conventions that secure the other basic liberties. It is *not* “perfectly obvious” that the economic liberties are not basic liberties, even if it is true. As Freeman argues, treating the economic liberties as basic is a fundamental commitment of an entire tradition of liberal thought:

I have discussed the main differences between the classical and high liberal traditions. While they both endorse personal liberties as fundamentally important, the classical tradition also gives priority to robust if not unqualified rights of private property in productive resources and other economic liberties, regarding them as of nearly equal significance with other basic liberties (2011, p. 52).

Conventionalists have avoided arguing against this substantive point by proceeding as though the legal instantiation of property makes it obvious that there is no independent or, as Breakey calls

it, *core* right to private property. Yet the denial of independent moral property rights finds no support from the trivial observation that *legal* property rights are legal conventions.

### 3.2.7 Implications of Exceptionalism

What are the implications of this diagnosis for the conceptual argument? First, a summary of the diagnosis itself. All the basic liberties are core rights or freedoms which are situated in legal conventions. Therefore, observing that property is a legal convention does not rule out a core right or freedom of economic liberty. Whether property rights are *mere* conventions (that is, there is no core right) is an open question, and not one which is addressed by the conventionalists. I think this diagnosis is basically right, as far as it goes. But how far does it go? Not, I will suggest, as far as critics of the conceptual argument seem to think.

Those critics proceed as though rescuing the possibility of a core right to private property defeats the conceptual argument. For example, with the conceptual argument out of the way (or so he thinks), Tomasi reiterates what he takes to be the populist credentials of the property constraint: “many Americans believe they have a right to the money they earn from working, and that proposals for taxation must therefore overcome that preexisting right” (2012, p. 61). He offers the elitist and out-of-touch demographics of professional philosophers as an explanation of why they erroneously “treat the settled convictions of ordinary working people that they have some sort of de facto claim on the wages they earn as meriting no more professional attention than the objectively false beliefs such people have in response to complex problems in probability theory” (2012, p. 215). Similarly, Breakey assumes he has vindicated the pretax baseline as “normatively significant but not conclusive” (2014, p. 600). This is so because the pretax baseline is, on his view, a *prima facie* property right that is not yet fully situated in a full

scheme of rights. Like Tomasi, he thinks ordinary people have understood this all along. Some academics have unfairly dismissed this ordinary understanding as confused “everyday libertarianism” but, according to Breakey, “[t]he folk are right and the theorists are wrong” (2014, p. 602). Gaus also thinks that pretax rights imply that the conceptual argument is “an error” (2010, p. 59).

I logically can have an entitlement to all my pretax income in the sense that taking away any of it must be publicly justified: since my pretax income is something I have a right to, any infringement of that right must be justified. Murphy and Nagel are certainly correct that some activities of the state are necessary for my property rights to exist: funds required for those activities are justified claims against my property. But that someone has a justified claim against some of my property does not show that I do not have “any kind of entitlement” to that part of my property.

In contrast, my view is that the conceptual argument does not turn on vindicating conventionalism. Perhaps counterintuitively, success of the conceptual argument is consistent with moral rights to property. The debate over the property constraint has nearly always been framed as between mere conventionalists on the one side arguing that “a natural right to property is unjustified” (2014, p. 574) and those who accept a natural or constitutional right to property on the other. I say ‘nearly’ because there are two exceptions of which I am aware. These two exceptions endorse *the independence thesis* which is the view that the soundness of the conceptual argument is independent of conventionalism. This might seem surprising, given Murphy and Nagel’s earlier cited claim that “[t]he conviction that determines our approach to all more specific questions is that there are no property rights antecedent to the tax structure” (2002, p. 74). Yet I suggest that the conceptual argument, which I take to be the central thesis of their book, is actually independent of that conviction.

The first of the two endorsements of the independence thesis comes from a short review of *The Myth of Ownership* by Eduardo Frajman and Joe Oppenheimer:

It is important to realise that Murphy and Nagel are making two related yet separate arguments. The more technical argument clarifies the role of taxes as an element of the overall system of property. This argument may be accepted at the same time as the authors' normative prescription that all property be viewed as conventional is rejected (2003, p. 385).

They seem to have recognized that the conceptual argument works without any commitment to conventionalism. As reviewed above, this is not at all how critics have responded to the conceptual argument. Instead, defenders of the property constraint assume that those who reject it are committed to a merely conventionalist account of property. This explains why Tomasi, Breakey, Gaus, and so on, all think a core right of property, or a commitment to economic liberty as basic, necessarily implies the property constraint. And, once again, this is not surprising given Murphy and Nagel's unnecessary insistence on a *merely* and *purely* conventional account of property, without ever explaining what those qualifiers mean.

The second endorsement of the independence thesis comes, not without irony, from Liam Murphy himself. His statements in 2005 emphasize how, in his view, the conceptual argument does not depend on conventionalism, though this seems to have been ignored by later critics. He writes that "*whatever the truth about the morality of property*, [emphasis added] we couldn't possibly have a right to the whole of our pretax income (Murphy, 2005, p. 124). This essentially reiterates his point that the conceptual argument is 'entirely general': "[w]hatever we believe we are entitled to as a matter of distributive justice *or property theory*, [emphasis added] or anything else, we cannot be entitled to the whole of our pretax income (Murphy, 2005, p. 124)."

Persistent focus solely on legal conventionalism has obscured the importance of the independence thesis in assessing the conceptual argument and, therefore, the property constraint. The result has been a kind of interpretive chaos. Consider Tomasi's claims that "Americans believe they have a right to *the money they earn* [emphasis added] from working" (2012, p. 61).

He apparently thinks this is contrary to Murphy and Nagel's view of things, yet they maintain that "[t]he money you earn under any system is yours because you have worked for it, but *it is a mistake to think that what you have really earned is your pretax income*, [emphasis added] some of which the government then comes and takes away from you" (2002, p. 63). Evidently all agree that people are entitled to what they earn. The fundamental question at stake is whether people earn their pretax income. The conceptual argument concludes that they do not, and the independence thesis implies that this can be true *even if* there is a powerful core right to private property.

To see this clearly, assume there is a core right to private property. Breakey (2014) and others see conventionalism as maintaining that "a natural right to property is unjustified" (p. 574). Therefore, suppose the economic liberties are, as the classical liberals think, nearly on a par with other basic liberties. The first premise of the conceptual argument is that property rights are legal rights. One way of interpreting that proposition is that there are no moral property rights independent of legal conventions. That is the interpretation promoted by Murphy and Nagel. However, since we assume a core right to private property, we reject that interpretation. What remains? As I indicated above, we retain the trivially true interpretation of the first premise: all *legal* property rights are legal rights. The crucial point here is that even with the trivially true premise the argument still goes through. The argument does not indicate what anyone *is* entitled to, but it picks out one distribution *they cannot be* entitled to: the legal pretax distribution. That, I take it, was the essential point all along:

1. [Legal] property rights are legal conventions.
  2. Legal conventions are created (in part) by taxation.
  3. Therefore, [legal] property rights are created (in part) by taxation.
  4. It is logically impossible for the creation (in part) of an institution X to provide an entitlement baseline independent of X for assessing the legitimacy of X.
-

5. Therefore, it is logically impossible for [legal] property rights that are the creation (in part) of the institution of taxation to provide an entitlement baseline independent of taxation for assessing the legitimacy of taxation.

This argument against entitlement to some legally described pretax distribution does not depend on any kind of deep conventionalism since it goes through even when we accept only the trivial claim that legal property rights are legal conventions. Therefore, the independence thesis is correct: whatever one thinks about property, there can be no entitlement to the pretax baseline. If there is a core right to private property, situating it cannot involve appeals to entitlement to some pretax distribution, since, if the conceptual argument is sound, there is no such entitlement. But, the soundness of that argument is independent of the conventionality of property with the exception of the trivially true claim which no one rejects: all legal property rights are legal conventions. As it turns out, the conceptual argument is consistent with even very strong economic liberty rights.

It is easy enough to see how the independence thesis has been overlooked. Murphy and Nagel insist that private property is merely conventional, and their main conclusion is that people are not entitled to some pretax distribution. But why would they insist on *mere* conventionalism if it is independent from their rejection of entitlement to any pretax baseline? That is a question for which I have no answer. However, confusion about the independence thesis does explain why critics have thought that undermining conventionalism rescues entitlement to pretax property rights when, in fact, it does no such thing. It also clarifies what is really at stake in the debate over the property constraint. For example, earlier I cited Tomasi's claim that "[w]ith respect to assertions about the legitimacy of restrictions on economic liberty, the claim that such liberties are legal conventions is trivial" (2012, p. 71). I think he is right about this. However, if the

conceptual argument stands (as I have argued it does), then taxation is not well understood as a restriction on economic liberty. After all, people are not entitled to some pretax distribution.

Another way of putting this same point is that the conceptual argument is consistent with economic liberty rights. Taxation is only a ‘restriction’ on those rights if they imply entitlement to *the precise (legally described) distribution of pretax income*. But, as the conceptual argument shows, that is incoherent. The conclusion holds even when the first premise is understood as strictly the trivial claim that all legal property rights are legal rights. Therefore, Tomasi has unjustifiably moved from claims defending economic liberties in general to a specific conclusion about pretax income as a baseline of entitlements. His defense of economic liberties is simply beside the point, since all the conceptual argument implies is that economic liberties, whatever they imply about legal property rights, cannot imply entitlement to all of pretax income.

### **3.2.8 Moral Property Constraints on Taxation**

One way to caricature the debate over the property constraint is this. There are natural rights theorists and mere conventionalists. Natural rights theorists think there are natural property rights. Therefore, people have a provisional right to their full market income, and government interference via taxation requires justification. Mere conventionalists think there are no natural property rights. Therefore, people are not entitled to their full market income, and what they are entitled to follows from taxation.

So far, I hope to have made two main points about the real substance of this debate. First, the ‘mere’ has been profoundly misleading: no one (including Murphy and Nagel) is really a mere conventionalist. Second, success of the conceptual argument (and so defeat of the property constraint) cuts across both the conventionalist and natural rights positions. Since the conceptual

argument does not depend on mere conventionalism, it is consistent with moral property rights that constrain taxation. Any debate ought to be about the substance of those constraints. The final point I hope to make is one of emphasis rather than argument. Critics and defenders of the property constraint alike emphasize the legitimacy (or injustice) of a *reduction* in after-tax resources as compared with the pretax baseline, even when there is no reason to favour that emphasis. A paper by Timothy Fowler (2015) helps illustrate all three of my points (which are, largely, the opposite of his own). I will outline the relevant points in his discussion, note which I reject, and then draw on his thought experiment to illustrate how they come together.

Fowler begins by putting natural rights theorists on one side and conventionalists on the other: “the proper distinction between natural rights theorists and conventional theorists is whether a property claim is an input into deciding which scheme of property laws to enforce or if it is purely an output of these laws” (2015, p. 350). For conventionalists, he maintains, property rights are ‘purely an output’ of law. This is the first point I reject; there are no mere conventionalists. Next, he draws the familiar implication: “[i]f one accepts a conventional view of property, it is not true that taxation can violate a property right claim” (2015, p. 351). He takes this to imply that “taxation, cannot, even in principle, violate property rights” (2015, p. 350). This is the second point I reject; taxation can violate property rights. Finally, he provides an example of Tribal Elders justly taking excess resources from workers. This illustrates my point about emphasis; taxation is not taking. To see how my three points bear on the discussion, I turn to Fowler’s thought experiment. For him, the central point of the example is to show that taxation cannot violate property rights. For me, the central point is to show that it can.

The thought experiment imagines a tribe that subsists on picking berries from a single large bush. The available berries are sufficient to sustain the tribe, but they are scarce. Fowler

describes various property schemes that might arise. Of interest is the final proposal, according to which gatherers can pick up to 200 berries each day but must contribute half of whatever they pick to a communal pot which feeds the unlucky. Fowler (p. 350) points out that for gatherers operating under this proposal “it is simply not true to say that people have any property claim at all to berries they picked after taking 100.” In general, the Tribal Elders in no way harm gatherers when taking any excess, as the gatherers lack even a *pro-tanto* claim to berries in excess of their allotment under the accepted scheme of property. The entitlement convention is what creates the property rights in any berries at all. This, he argues, is “analogous to state taxation on the conventionalist view of property” (Fowler, 2015, p. 350).

Again, my first point is that there are no mere conventionalists. Notice that Fowler assumes the property convention justifies any resulting entitlement. If, as a matter of fact, some scheme of property is practiced by the tribe, then it dictates property rights. However, even arch-conventionalists Murphy and Nagel rightly insist that entitlement depends on a system that is *fair* and *legitimate*: “income gives rise to clear moral entitlement only if the system under which it is earned, including taxes, is fair” (p. 74) and “what [pretax income] legally entitles you to is morally legitimate only by virtue of the legitimacy of the system” (p. 75). Suppose the tribe’s property convention is (for whatever reason) unfair or illegitimate. Then what the convention defines as “excess” is not necessarily something over which the Elders or the rest of the community have a justified entitlement. The gatherer may be morally entitled to either more or less than the convention allows.

The key point is that what is grounding the gatherers’ entitlement is *not* the convention *per se*, but rather the normative fact that the convention is *fair* or *legitimate* (if it is). Fowler will respond (as he indicates in footnote 3) that this only proves that there are good reasons to reject

“Murphy and Nagel’s purely conventionalist view.” As he argues, appeals to the fairness and legitimacy of the convention are inconsistent with pure conventionalism. Yet these appeals to the necessity of a fair and legitimate system for creating property rights are Murphy and Nagel’s own words! So what Fowler really demonstrates is, as I have emphasized, that Murphy and Nagel are *not* ‘pure’ or mere conventionalists at all. Mere conventionalism is a red herring in the debate. All accept that powerful normative reasons must guide any system of private property.

My second main point then follows from the first: conventionalism does not rule out moral property rights as constraints on taxation. As we have seen, since there are no mere conventionalists, a convention can only generate entitlements if it is legitimate. The mere fact that it *is* a property convention does not mean it creates property rights. Even Murphy and Nagel sensibly insist that a wide range of political values may be brought to bear in assessing the legitimacy of any convention. Might these considerations include economic liberties or a core right to private property? The mere fact that property is realized in law does not rule them out. Such claims must be proposed and defended like any others. If conventionalists want to deny that economic liberties are among those reasons, then they must do so on substantive grounds, rather than claiming those liberties are ruled out by adherence to a ‘pure’ or mere conventionalism.

My final point is simply to highlight the regrettable emphasis on reductions from some pretax baseline. Arguments about the property constraint are naturally framed in terms of complaints about taxation taking resources. The system of taxation appears to leave people worse off than they would be if they were entitled to all of their pretax income. But, it is said, people only appear worse off, since they had no entitlement to that amount in the first place. Thus the arguments point out that people are not entitled to keep *all* of their pretax income, or the *whole* of their pretax income, and so on. Yet really the conceptual argument shows that the pretax

income figure is morally arbitrary. It simply does not influence what people are entitled to as a matter of justice. But this provides no reason at all to think that what people are entitled to is *less* than the pretax figure. In many cases it may be more.

My own view is that people are often entitled to more resources than they have. And taxation often interferes with those entitlements. I doubt there should be any taxes on income levels below, at minimum, a living wage. People working full-time in a wealthy nation should not labour under an economic system (including taxation) that leaves their after-tax resources *below* what is required for a reasonable standard of living. When the tax system leaves such people worse off than it should, then taxation is ‘interfering’ with their moral right to resources—it is failing to secure their property rights; indeed, it is actively violating them. It is in this sense that moral property rights can pose a constraint on taxation. The correct standard of justice may, and in my view does, require allocating many people far more resources than they receive under the current institutional arrangement. But my appeal to moral property rights that constrain taxation makes no reference to any phantasmal ‘pretax’ entitlements. I do not think people are morally entitled to some pretax figure derived from a calculation made by the legal system, but nor do I assume their entitlements are *less* than that figure. In short, I find it unfortunate that conventionalists implicitly accept the framing of the conceptual argument as proving that ‘you are not entitled to *that much*’ when they could better frame it as ‘you are not entitled to *that number*.’ As Sampford and Wood (1991, p. 183) put it, we should only ask “how much property should this person have?”

### 3.3 Chapter 3 Conclusion

Over the past half-century, a small number of theorists have observed that there seems to be something incoherent about taking a preliminary legal description of property holdings as a fixed baseline of profound moral salience against which the interferences of taxation can be measured. Their target is captured by Duff's casual identification of the correct baseline for assessing vertical and horizontal equity: "the appropriate reference point, given that the income tax is designed to be levied based on taxpayers' income, is income" (2018). That this baseline can be offered with no justification other than that it is "appropriate" is striking. The most prominent argument that has been brought against it maintains that property is nothing more than a legal convention, and so it makes no sense to characterize a pretax description of property as a *moral constraint* on the very legal system that is its genesis.

Critics have not been persuaded by this line of argument. They point out that property is not merely a legal convention. Property rights exist outside the law, and in social contexts without law. Even when property laws exist, practices of property often diverge from and conflict with those laws. It is of course true, they allow, that the moral force of property is realized through particular legal conventions, but this seems not different from all basic rights, which are similarly legally realized and contextualized. With this more accurate conception of property in hand, critics have rejected the conceptual argument as depending on an implausibly conventionalist account of property.

I argue that the conceptual argument is independent from an implausibly conventionalist account of property, and indeed from any theory of property at all. I note that the conceptual argument only depends on characterizing preliminary descriptions of holdings (such as pretax income) as legal constructs. Whatever one thinks about property, the legal calculation of pretax

income is undeniably a product of the legal system. This observation is *independent* from any particular theory of property. Even if the goal of property institutions is to realize fundamental rights to economic liberty, it is still the case that the legal description of pretax holdings is, after all, a legal description. Since that description depends on a system of which taxation is an inextricable part, the conceptual argument is untouched by the critics' rejection of conventionalism.

## Chapter 4: The Baseline Problem in Tax Efficiency

### 4.1 Chapter 4 Introduction

The purpose of this chapter is to reflect on problematic baselines latent in simple expressions of two economic concepts that are significant in a great deal of thinking about taxation. I am not an economist, and therefore my intention is not to critique the use or value of these concepts in the technical work of economists. Yet most people who think and write about taxation are also not economists, even though academic economics dominates *theoretical* work on taxation. Lawyers, policy analysts, philosophers, political scientists, and many more non-economists all discuss taxation, and they inevitably draw on economic concepts when doing so. This naturally involves exporting certain economic ideas from the model environments in which they were developed to practical environments which deviate significantly from model conditions. This exporting of concepts from model environments to practical ones is meant to assist in thinking about taxation, but it involves some risk. The risk I identify here is a failure to notice the way in which both of these concepts are strictly comparative, and so to pay insufficient attention to their comparative baselines. Further, focus on these baselines tends to undermine the economic insights they are supposed to deliver.

In this chapter, the two economic concepts I consider are “Pareto improvement” and “deadweight loss.” These are familiar concepts to most people, and they have notable influence over public policy in general and tax policy in particular. This is because it is widely agreed that taxes should be as efficient as possible, and these concepts are foundational to thinking about integrating private and government action to maximize efficiency. Yet for all their influence, I think the utility of both concepts in evaluating tax policy is more limited than is commonly

recognized. I argue that their limitation is related, in both cases arising from reliance on problematic baseline thinking.

The first section considers the Pareto principle. In brief, Pareto improvements are changes which at least some people prefer, and no one disprefers. If some change in tax rates makes some people better off and no one worse off compared with distribution **A**, then it seems like a straightforwardly efficient and desirable improvement over **A**. My aim in this section is not to cast doubt on that compelling claim. Instead, I want to draw attention to the way in which such comparisons risk privileging a status quo baseline (**SQ**) distribution without justification. The risk plays out in the following way. We notice that **A** is a Pareto-improvement over **SQ**, while **B** is not. It is therefore inferred that the Pareto principle (which recommends Pareto-improvements) provides a reason to prefer **A** over **B**. But this is false, as the Pareto principle only provides a reason to prefer **A** over **SQ**. This simple error has surprising reach, as I illustrate through Joseph Heath's recent book on the topic.

The second section considers deadweight loss. In the tax context, deadweight loss refers to unrealized economic gains caused by taxation distorting market behaviour. By the late 1980s, the Musgraves reported "growing attention...in tax theory" to "efficiency costs or deadweight losses which do not arise in a market for private goods" (1989, p. 8). Today, this attention persists, especially with the rise of optimal tax theory. Like the notion of a Pareto-improvement, the notion of deadweight loss is strictly comparative. It is a measure of net loss between a taxed and untaxed market. As Fleurbaey describes it, economists can "compute the deadweight loss due to taxation, as compared to the laissez-faire allocation" because "laissez-faire is in the set of feasible allocations ('no tax' is a particular tax policy)" (2018, p. 42). However, if the property, contract, tort, and other legal rules that help constitute the market fundamentally rely on a legal

system, then it is not clear how “no tax” is a coherent policy choice. In practice, the very economic power at stake in tax policy is legally defined property rights. As emphasized in chapters one and two, those legal rights are created by a system that already relies on taxation. Therefore, it seems confused to hold the distributive outputs of that system constant while its antecedent conditions are assumed away. It implies that without taxation the market surplus would be much larger than it is, but in fact there would be nothing like a modern market at all without taxation (no corporations, no legal property rights, no courts, no state enforcement, and so on).

## **4.2 The Pareto Principle**

### **4.2.1 Introduction**

One common way to understand efficiency in tax analysis is Pareto efficiency. In this section, I introduce the relevant cluster of Pareto concepts. I observe that Pareto efficiency holds a special place in public policy in general, and in tax policy in particular. Some defenders of the Pareto principle argue that it deserves this place because it is value neutral in at least some important sense. I consider Joseph Heath’s recent defense of this claim but conclude that any apparent neutrality is illusory. My argument is that the appearance of neutrality follows from implicitly and arbitrarily privileging the status quo as a baseline in Pareto comparisons. This unjustified privilege normally goes unnoticed due to a conflation of two claims: (1) the mostly uncontroversial claim that Pareto-superior allocations are better than their Pareto-inferior alternatives and (2) the much more controversial claim that allocations Pareto-superior to the status quo are better than their Pareto-incomparable alternatives. Taken as a whole, my criticism

suggests that applying the Pareto principle to public policy choices faces a serious and overlooked baseline problem.

#### 4.2.2 Pareto Concepts

The first concept in the Pareto cluster is *Pareto efficiency*. An intuitive (though, I will soon suggest, somewhat misleading) definition is that a distribution of resources is Pareto efficient if and only if no possible reallocation can make anyone better off without making someone worse off. In tax policy, this concept is simply applied to tax schedules: “The allocation generated by some tax schedule is ... Pareto efficient if there is no other tax schedule that induces a resource feasible allocation where nobody is worse off and some workers are strictly better off” (Werning, 2007, p. 3). The obvious implication of this definition is that whenever a distribution is not Pareto efficient, reallocation is possible which makes some people better off without making anyone worse off. Put another way, a Pareto inefficient allocation makes some people worse off without benefitting anyone. When an allocation is not Pareto efficient, reallocations that make some better off while making none worse off are called *Pareto improvements*. If some are better off and none are worse off in  $x$  than  $y$ , then  $x$  is *Pareto-superior* to  $y$  and  $y$  is *Pareto-inferior* to  $x$ . Crucially for my argument, if some are better off and some are worse off in  $x$  than  $y$ , then  $x$  is *Pareto-incomparable* with  $y$ . Finally, there is the *Pareto principle*, which is the normative commitment that Pareto-superior allocations are better than Pareto-inferior allocations. In plain English, then, the Pareto principle holds that “If a change benefits everyone, then it is a change for the better” and also, “if a change benefits some but harms no one, then it is a change for the better” (Fisher & McClennen, 2011, p. 7).

The intuitive appeal of the Pareto principle is widely recognized. Indeed, it is often accepted as self-evident. When considering a shift from a Pareto-inferior to a Pareto-superior state of affairs, Fisher and McClennan ask,

Why would anyone object to such a shift insofar as it is *ceteris paribus* a ‘change for the better’? Such considerations can be thought of as instances of the more general claim that if someone can benefit from something without harming anyone else, then we ought to let such a state of affairs be realised. The self-evidence of which stares us in the face (Fisher & McClennan, 2011, p. 7).

I should note that they are not ultimately satisfied with the epistemic stability of self-evidence, and so they go on to ground the Pareto principle in individual rationality. But their “stares us in the face” account of self-evidence does capture the intuitive power of preferring allocations that make some better off while harming none. As James Griffin earlier wrote, “It is hard to see how anyone could resist such a principle” (1986, p. 147).

I previously qualified my introductory definition of Pareto efficiency as “somewhat misleading” because in Pareto analysis “better off” and “worse off” are understood in a specific way: as individuals’ preference rankings over states of affairs. Yet it is perfectly sensible to say that someone prefers  $x$  over  $y$  but that in fact that person is better off under  $y$ . Critics of preference utilitarianism, for example, have often pointed out that “some preferences are misinformed, crazy, horrendous, or trivial” (Sinnott-Armstrong, 2023). Perhaps the least controversial here is misinformation: “I might prefer to drink the liquid in a glass because I think that it is beer, though it really is strong acid” (Sinnott-Armstrong, 2023). Clearly, my misinformed preference diverges from what makes me better off. Yet Pareto analysis does not track this divergence of well-being or “better offness” from preference satisfaction. That is, Pareto analysis relies on preference rankings which do not necessarily correspond to a person being better or worse off. Therefore, “It is more accurate to say that  $R$  is a ‘Pareto improvement’

over S if nobody prefers S to R and somebody prefers R to S” (Hausman and McPherson, 2016, p. 65). This clarification from “better off” to preference satisfaction can be generalized to the full cluster of Pareto concepts, including the Pareto principle. Therefore, to accurately capture this dimension of Pareto analysis, my use of the principle follows Amartya Sen’s formulation:

- (a) If at least one individual strictly prefers  $x$  to  $y$ , and every individual regards  $x$  to be at least as good as  $y$ , then the society should prefer  $x$  to  $y$ .

### **4.2.3 Neutrality in Public Policy**

The Pareto principle is intuitively attractive, and it is difficult to see how anyone could reasonably object to changing a state of affairs to an alternative that everyone agrees is at least as good, and at least some think is better. While people can be wrong about the value of some state of affairs, the world of public policy understandably prefers to avoid engaging in such value judgments. Still, it is important to recognize that the Pareto principle is not strictly value neutral. After all, it ranks allocations as better or worse than others. Moreover, the grounds on which it delivers such rankings inevitably include normative assumptions. For instance, Hilary Putnam argues that accepting an individual preference interpretation of Pareto efficiency ultimately relies on “the underlying value judgment that every agent’s right to maximize his or her utility is as important as every other’s” and so it “isn’t a value neutral criterion of ‘optimality’ at all” (Putnam, 2004, p. 56). I would add that the individual preference rankings that serve as inputs into the principle are themselves products of a particular state of affairs (namely, the status quo), and there may be reasons not to accept them as authoritative. For example, someone may prefer  $x$  over  $y$  under conditions of severe economic inequality, but that preference may be reversed under different economic conditions. As Bryan observes, market logic takes interests for granted,

when “the essence of the interest itself is bound up with its institutional manifestation” (2000, p. 7).

Though the Pareto principle is not strictly value neutral, it still embodies a certain kind of neutrality, and one that many have found compelling. In his recent book, Joseph Heath (2020) argues that the Pareto principle deserves its special place in guiding public policy because it is able to rank states of affairs without appealing to comprehensive and controversial theories of the good. It judges allocations as better and worse, but it does so with full deference to the preference rankings of individuals. The state must make choices among policy alternatives, but citizens hold different conceptions of the good, and state attempts to coercively impose the whole truth on the public have been disastrous. Thus, he argues, modern liberalism emerged as a “a second-best solution to the problem of political order” (p. 136). Instead of the state determining that  $x$  allocation is better than  $y$  in terms of some comprehensive theory of the good, it aims to realize allocations that everyone agrees are improvements on the status quo. Heath argues that this kind of neutrality has a proven history of promoting cooperation and minimizing social conflict against a background of value pluralism (p. 186).

Heath’s central thought is that if  $x$  is “deemed better for all, each from his or her own perspective” (p. 193) then the state can pursue  $x$  without violating its commitment to liberal neutrality. And it is of course a key feature of Pareto improvements that they make some people better off and none worse off, *by their own lights* (p. 84). Thus, on the one hand, liberal neutrality opposes “policies that can only be justified through appeal to particular controversial values” (192). And on the other hand, the Pareto principle ranks “some states of affairs as *better than others*” while taking no stand on particular controversial values (p. 144). Pareto efficiency is therefore celebrated by Heath as “the paradigm instance of a normative principle that the state

can act on without presupposing the correctness of any first-order conception of the good” (p. 193). In summary, Pareto improvements can be pursued largely without controversy since they harm no one (p. 144). They are better for everyone, easy to implement, an all-around win-win.

Given the special features described, Heath thinks the Pareto principle should play a major—indeed the primary—role in guiding public policy. Government pursuit of uncontroversial improvements is the normative cornerstone of not only taxation, but the modern welfare state in general. Of the three “liberal” principles he identifies, “efficiency is without doubt the most important. (p. 346)” It is the “central principle” (p. 86) of the executive branch, and the expansion of the welfare state was guided “above all [by] the commitment to efficiency” (p. 92). And efficiency--not equality or aggregate utility--is the dominant rationale for cost-benefit analysis, the major spending programs (health, pensions, education), regulatory interventions, and the social safety net (p. 149). Accordingly, the role of the Pareto principle in public policy decisions is justified and should even be expanded. It is not normatively inert, but it competently guides policy without presupposing any particular conception of the good. It is, we might say, a thin, uncontroversial moral commitment appropriate to modern liberal democratic contexts.

#### **4.2.4 An Example Choice**

What follows is a simple case to facilitate assessing the Pareto principle as a relatively neutral normative commitment that can guide “both policy choice and implementation” (p. 84). Suppose the Ministry of Finance is advising on a choice between three feasible tax schedules. For simplicity, I will characterize the affected population as just three individuals and I will assume their preference rankings align with their resource allocation.

Table 2: A Choice Among Three Tax Schedules

Tax Schedule		Resource Allocation		
		Person 1	Person 2	Person 3
<b>SQ</b>	Status quo	2	5	8
<b>DT</b>	Decrease tax	2	5	10
<b>IT</b>	Increase tax	4	6	7

Heathian defenders of the Pareto principle see this as an easy choice. The Ministry of Finance analysts look at the choice set and notice that **IT** makes some people better off, and that it is a more equal distribution than **SQ**. However, **IT** makes **person 3** worse off and so will be controversial to that extent. In contrast, everyone agrees that **DT** is at least as good as **SQ**, and some strictly prefer **DT** to **SQ**. Presumably a number of considerations will come to bear on the final decision among these tax schedules, but the Pareto principle is ostensibly a powerful justification in favour of **DT**. “Other choices are controversial,” one analyst might say, “but **DT** is a Pareto-improvement.” “That is, everyone agrees that **DT** is equal to or better than what they’ve got.” The state selecting **IT** might require appeal to a more comprehensive theory of the good, while choosing **DT** is more neutral—an uncontroversial improvement.

#### 4.2.5 Comparative Baselines

I think this way of framing public policy choices is mistaken, and that attention to its arbitrary privileging of the **SQ** baseline reveals how it goes wrong. The illusion of neutrality—that **DT** is much less controversial than **IT**—is caused in part by describing states of affairs as “Pareto improvements,” which tends to obscure the fact that the Pareto principle is “strictly comparative” (Wonnell, 2001, p. 692). In general, describing **x** as a Pareto improvement amounts to saying that every individual is indifferent to or prefers **x**. Yet such a description is empty

unless accompanied by some alternative over which **x** is preferred. It is akin to saying that **x** is *better* (as Heath and others often do) without saying *what* it is better than. Note that formal definitions of the Pareto principle, such as Sen’s referenced above, avoid obscuring the comparative baseline:

- (a) If at least one individual strictly prefers **x** to **y**, and every individual regards **x** to be at least as good as **y**, then the society should prefer **x** to **y**.

Sen recommends that, when (a) is fulfilled, we say “**x** is Pareto-wise better than **y**” (2017, p. 21). Similarly, I said above that **x** is Pareto-superior to **y**. This kind of formulation highlights the fact that a Paretian analysis necessarily depends on specifying the distributive structure that forms the comparative baseline **y**. But critically, no particular baseline is “imposed by the analysis itself” and above all “[t]he background distribution of rights assumed to be constant in a Paretian analysis need not be the status quo” (Cooter, 1989, p. 822). Different rights structures will yield different sets of improvements relative to those rights structures (Samuels, 1981, p. 125). Therefore, verdicts of the Pareto principle are perhaps surprisingly limited. The principle “says merely that an efficient change from the status quo is better than the status quo, and that an efficient change from other baselines is better than those baselines” (Wonnell, 2001, p. 697).

Bearing in mind the strictly comparative nature of the Pareto principle, we can return to the choice among tax schedules. Note that the options have the following Pareto comparisons.

*Table 3: A Choice Among Three Tax Schedules (repeated)*

Tax Schedule		Resource Allocation		
		Person 1	Person 2	Person 3
<b>SQ</b>	Status quo	2	5	8
<b>DT</b>	Decrease tax	2	5	10
<b>IT</b>	Increase tax	4	6	7

Table 4: Pareto Comparisons of Three Tax Schedules

Tax Schedule		Pareto comparisons
<b>SQ</b>	Status quo	Pareto-incomparable to <b>IT</b> Pareto-inferior to <b>DT</b>
<b>DT</b>	Decrease tax	Pareto-superior to <b>SQ</b> Pareto-incomparable to <b>IT</b>
<b>IT</b>	Increase tax	Pareto-incomparable to <b>SQ</b> Pareto-incomparable to <b>DT</b>

Given these comparisons, what guidance can the Pareto principle provide? Certainly, it entails a social preference for **DT** over **SQ**. Assuming the Pareto verdict is the dominant consideration in this choice, the Pareto principle rules out **SQ**. It holds that, given a choice between **DT** and **SQ**, society should opt for **DT**. Yet this is *all* the principle can recommend. Unfortunately, this is not the choice faced by the Ministry of Finance. Instead, the choice is among three options, some of which are Pareto-incomparable. If we try and apply the principle to the remaining choices, we see that choosing **DT** over **IT** is not recommended by the Pareto principle. Likewise, choosing **IT** over **DT** is not a violation of the Pareto principle. This is because **Person 3** prefers **DT** to **IT**, while **Person 1** and **Person 2** prefer **IT** to **DT**. As a result, **DT** and **IT** are Pareto-incomparable, and so the Pareto principle is simply silent as between **DT** and **IT**. Certainly the principle has provided some limited guidance: do not choose **SQ** when **DT** is available. But in a world of multiple competing alternatives, that does not support choosing **DT** over Pareto-incomparable alternatives like **IT**. Thus, the Pareto principle cannot recommend choosing **DT** in this example.

#### 4.2.6 Uncontroversial Conflation

Nevertheless, Heath contends that in such cases the Pareto principle *does* recommend **DT**. Why? Because **DT** is “a Pareto improvement” while **IT** is not. Notice how this conflates the Pareto principle giving a reason to prefer **DT** over **SQ** (which it does) with the principle giving a reason to prefer to **DT** over **IT** (which it does not). Again, the Pareto principle only gives a reason to prefer Pareto-superior over Pareto-inferior allocations.

The conflation here is subtle, but it is important to distinguish two ways in which the Pareto principle could be thought to provide reasons in choices like this example. We might think of these as first and second order reasons. The Pareto principle provides a first order reason when it ranks one allocation above another, as it ranks **DT** superior to **SQ**. But one might also think that *the fact that* an allocation is so ranked is a reason to prefer it. This is a kind of second order reason. In our example, we know that the Pareto principle is silent as between **DT** and **IT**. Yet we might think *the fact that* **DT** is a Pareto improvement over **SQ** is a reason in its favour when choosing between **DT** and **IT**. In other words, **DT** has something going for it—it is a Pareto improvement over **SQ**—that **IT** does not. But notice that this justification invokes an entirely new normative commitment; namely, some kind of preference for allocations that are Pareto-improvements over **SQ**. Such a commitment is decidedly not part of the Pareto principle itself, and—crucially—does not enjoy the same kind of neutrality or universal support. This is clear when we specify (rather than implicitly assume) a principle that licenses a normative preference for “Pareto-improvements over **SQ**.” Here are two examples.

**(1) The *SQ* constraint:**

When choosing policies, the set of alternatives is constrained such that it includes only **SQ** and options that are Pareto-superior to **SQ**.

**(2) The SQ privilege:**

When choosing policies, a policy that is Pareto-superior to **SQ** is better than any alternative that is not Pareto-superior to **SQ**.

Whatever we think about these principles, they are obviously not the Pareto principle, and nor would they receive its widespread endorsement. Yet something like these is necessary to privilege **DT** over **IT** in the example case. Implicitly accepting such commitments is a common “habit of mind among economists” (Cooter, 1989, p. 822) but utterly at odds with a commitment to avoiding first-order value judgments. This is because, while **the SQ constraint** does justify choosing **DT** over **IT**, the constraint itself stands in need of justification. So building the **SQ** baseline into the analysis does deliver Heath’s conclusion that Pareto-improvements over **SQ** are justified, but only at the cost of the very neutrality that is supposed to justify the Pareto principle’s role in public policy in the first place. This should make explicit what Heath failed to notice—**SQ** functions as a privileged but undefended baseline in his analysis.

To emphasize the undefended privilege afforded **SQ** in the analysis, we need only imagine a fourth choice in our example case which is Pareto-inferior to **IT** but Pareto-incomparable to **DT** and **SQ**. Call this choice equal tax, or **ET**.

*Table 5: A Choice Among Four Tax Schedules*

Tax Schedule		Resource Allocation		
		Person 1	Person 2	Person 3
<b>SQ</b>	Status quo	2	5	8
<b>DT</b>	Decrease tax	2	5	10
<b>IT</b>	Increase tax	4	6	7
<b>ET</b>	Equal tax	3	6	7

Table 6: Pareto Comparisons of Four Tax Schedules

Tax Schedule		Pareto comparisons
<b>SQ</b>	Status quo	Pareto-incomparable to <b>IT</b> Pareto-inferior to <b>DT</b> Pareto-incomparable to <b>ET</b>
<b>DT</b>	Decrease tax	Pareto-superior to <b>SQ</b> Pareto-incomparable to <b>IT</b> Pareto-incomparable to <b>ET</b>
<b>IT</b>	Increase tax	Pareto-incomparable to <b>SQ</b> Pareto-incomparable to <b>DT</b> Pareto-superior to <b>ET</b>
<b>ET</b>	Equal Tax	Pareto-incomparable to <b>SQ</b> Pareto-incomparable to <b>DT</b> Pareto-inferior to <b>IT</b>

I earlier imagined a policy analyst citing the Pareto principle as a reason to choose **DT**, since the analyst claimed **DT** is a Pareto-improvement. But in this example, **IT** is also a Pareto-improvement. While **DT** is a Pareto-improvement over **SQ**, **IT** is a Pareto-improvement over **ET**. Therefore, we can imagine a different analyst preferring **IT**, also noting that **IT** is a Pareto-improvement. The key observation here is that this disagreement cannot be settled by the Pareto principle itself. The Pareto principle gives no reason to prefer Pareto-improvements over **SQ** more than Pareto-improvements over other baselines. Choosing **DT** implicitly relies on something like the example principles outlined above. For instance, the first analyst may

implicitly accept **the SQ privilege**, affording special significance to Pareto-improvements over **SQ**. But again, this is an additional normative commitment; it is not part of the Pareto-principle.

It is worth emphasizing that the provided example is not an edge case designed to undermine Pareto comparisons in certain unusual circumstances. Instead, the point is practically universal in public policy. It characterizes any case in which the set of alternatives includes a mix of “win-win” and “win-lose” options, relative to the status quo. These are exactly the cases Heath has in mind, as he notes that “win-win” outcomes minimize complaints, since there is no “‘natural’ constituency to press for [their] elimination” (2020, p. 177). Yet this is only true in the unlikely case in which the *only alternative* is “lose-lose.” Consider again the example above—is there no natural constituency to press against **DT**? If the only alternative is **SQ** then indeed there is no one to complain. But there is an additional alternative, and so the basis of complaint is clear: “True, I may be equal or better off in **DT** than **SQ**, but why is **SQ** the relevant baseline when under feasible alternative **IT** I am much better off?” Important policy choices are of course *not* between strictly “win-win” and “lose-lose” options. They are always among a set of alternatives that include allocations making some people better off and some worse off. The Pareto principle simply does not recommend pursuing Pareto improvements over the status quo instead of these Pareto-incomparable alternatives.

Heath makes a recent and sophisticated defense of Pareto neutrality. But he has not vindicated his central claim that efficiency “is the normative commitment that the civil service can most assertively promote without violating political neutrality” (p. 346). Instead, the Pareto principle manifests a deep tension between neutrality and efficacy. Free of first-order value claims defending a particular comparative baseline, its verdicts are limited. All it can recommend is not to choose *y* when *x* is Pareto-superior to *y*. The civil service presumably faces few such

choices. In contrast, defending a specific baseline empowers the Pareto principle to deliver clear verdicts in a wide variety of cases. The downside, of course, is that the illusion is dispelled and not even a veneer of neutrality remains.

#### **4.2.7 Pareto Principle Conclusion**

According to some, the Pareto principle promises to guide public policy towards outcomes that no one can reasonably reject. Across a sea of controversy, the Pareto principle identifies “win-win” solutions that the civil service should be quick to implement. Heath imagines cases in which, for example, everyone prefers (or is indifferent to) tax reform  $x$ , while preferences for tax reform  $y$  are mixed. He sees this as a powerful argument in favour of  $x$ . Yet the notion of a “preference for reform  $x$ ” fails to specify any comparative baseline. Everyone prefers reform  $x$  *to what?* The implicit baseline is inevitably the status quo. But the fact that everyone prefers  $x$  to the status quo, while the same is not true of  $y$ , may be *no reason at all to prefer  $x$  to  $y$* . For my argument, it only matters that, if it is a reason, it has nothing to do with the Pareto principle.

### **4.3 Deadweight Loss**

#### **4.3.1 Introduction**

In the tax context, deadweight loss refers to unrealized economic gains caused by taxation distorting market behaviour. It is “the total inefficiency produced by tax distortions, measured in terms of an equivalent loss in money (Fleurbaey, 2018, p. 42).” This “loss” is also called excess burden, referring to “the burden on individuals in addition to the revenue collected”

(Econometrics, p.1). That is, while taxes direct some market surplus to government, deadweight loss refers to the *net reduction* in surplus.

Economists use the term “deadweight loss” as a synonym for the distortions in economic decision-making caused by taxes. Thus, the “deadweight loss” of a tax is the welfare cost to society of the tax, beyond the actual cash payments made to the government, that are attributable to hours not worked, investments not made, or the like, because of the tax burden that would be imposed on those affirmative acts. No practical tax, whether a consumption tax or an income tax, can avoid the deadweight loss associated with distorting taxpayers’ preferences... (JCT, 2008, p. 54).

Accordingly, some deadweight loss is generally seen as an inevitable feature of taxation.<sup>51</sup>

Taxes create a deadweight or net loss to society. The greater the deadweight loss, the greater the economic inefficiency (Elkins, 2006, p. 47).

One obvious goal of tax policy, then, is to minimize deadweight loss as much as possible, given other constraints on tax design. There is a large literature aimed at achieving this goal, arguing that this or that tax proposal minimizes or worsens deadweight loss.

In choosing among alternative policy instruments, allowance must be made for resulting deadweight losses or efficiency costs, i.e., costs which arise as consumer or producer choices are interfered with (Musgrave & Musgrave, 1989, p. 11).

For example, many studies claim that taxes on capital income, savings, or other bases create more or less deadweight loss compared with alternatives, such as consumption (Fried, 1992, p. 961-962). Other studies try and estimate the deadweight loss of various tax measures, claiming for example that “efficiency costs or deadweight losses were estimated to amount to 20 per cent or more of [U.S.] income tax revenue” (Head, 1993, p. 30). Whatever one thinks about these specific findings, there is a general principle in public finance that “the deadweight loss inherent in alternative taxes should be considered in constructing a good tax system (Musgrave & Musgrave, 1989, p. 293).”

---

<sup>51</sup> In theory there are impractical head taxes (also poll tax or capitation) which are intended to be non-distortionary.

Analyses of deadweight loss have also contributed to common assumptions of tax policy. For example, tax models often show large distortionary effects at high marginal rates: “deadweight loss rises with the marginal tax rate” (Musgrave, 1990, p. 120). Indeed, a declining (or even zero) tax rate on high earners was an early finding of optimal tax models (Mankiw, 2009, p. 7). This contributed to the view that taxation must grapple with an inevitable conflict between equity and efficiency. That is, it is more efficient, but less equitable, to reduce tax on high earners. In 2008, the Joint Committee on Taxation accepted this finding, noting that,

In general, as marginal tax rates become higher, the deadweight losses (the distorting effects of taxes) imposed on society increase, because the higher marginal rates impel stronger taxpayer reactions. As a result, in designing a tax system or a tax expenditure there frequently is a direct tradeoff between equity and efficiency goals.

Like Pareto efficiency, the idea of deadweight loss appears simple and intuitively compelling. One need only compare the market surplus “post tax” to the market surplus in some free market equilibrium. These amounts differ because taxes drive a wedge between supply and demand, and the size of that wedge represents deadweight loss, or value foregone due to a suboptimal quantity of transactions. In short, without tax a certain number of transactions will occur. Once tax is imposed, the quantity of transactions will be reduced. The post-tax surplus will be distributed among consumers, producers, and the government, but some value is inevitably forfeited by this decline in the quantity of transactions. That foregone value is simply missed rather than redistributed, and so is a significant “cost” of the tax. This is a regrettable though at times necessary loss to society. All that remains is to minimize it as much as possible, while balancing it with other considerations such as equity.

Tax interferes with economic decisions and distorts efficient choice. This distortion is burdensome to you, while being of no help to the Treasury. Efficient policy should therefore minimize this burden, referred to variously as excess burden, deadweight loss, or efficiency cost” (Musgrave & Musgrave, 1989, p. 279).

However, also like Pareto efficiency, deadweight loss is a comparative concept. To see how this comparison works, as well as to facilitate my criticism of this concept in practical contexts, I refer to a simplified example of taxing the oil market, borrowed (with modification) from *Principles of Microeconomics* (Hutchison, 2017).

### **4.3.2 Taxing the Oil Market**

We begin with a tax-free market in oil: 4,000,000 gallons sold at \$4 per gallon. The government wants to levy a tax on oil at \$3 per gallon. This tax can be levied on either producers or consumers. If the government decides to levy the tax on producers, the marginal cost of production is increased by \$3 per gallon at every quantity level. This increase shifts the supply curve up, and a new equilibrium is created.




Figure 1: Oil Market with Tax on Producers

Given this \$3 per gallon increase in the cost of production, the supply curve now intersects with demand at \$5 per gallon, with 2,000,000 gallons sold. Demand has decreased due to the higher cost. Note that producers only receive \$2 per gallon, since \$3 per gallon is paid in tax.

Also note that (in the model) the economic incidence of the tax is the same regardless of whether it is legally levied on consumers or producers. To see this, suppose instead that the government added a \$3 per gallon tax on consumers at the register. This increase of \$3 per gallon at every quantity level shifts the demand curve down. Just as when the tax was levied on producers, the new equilibrium price is \$5 per gallon, with 2,000,000 gallons sold.

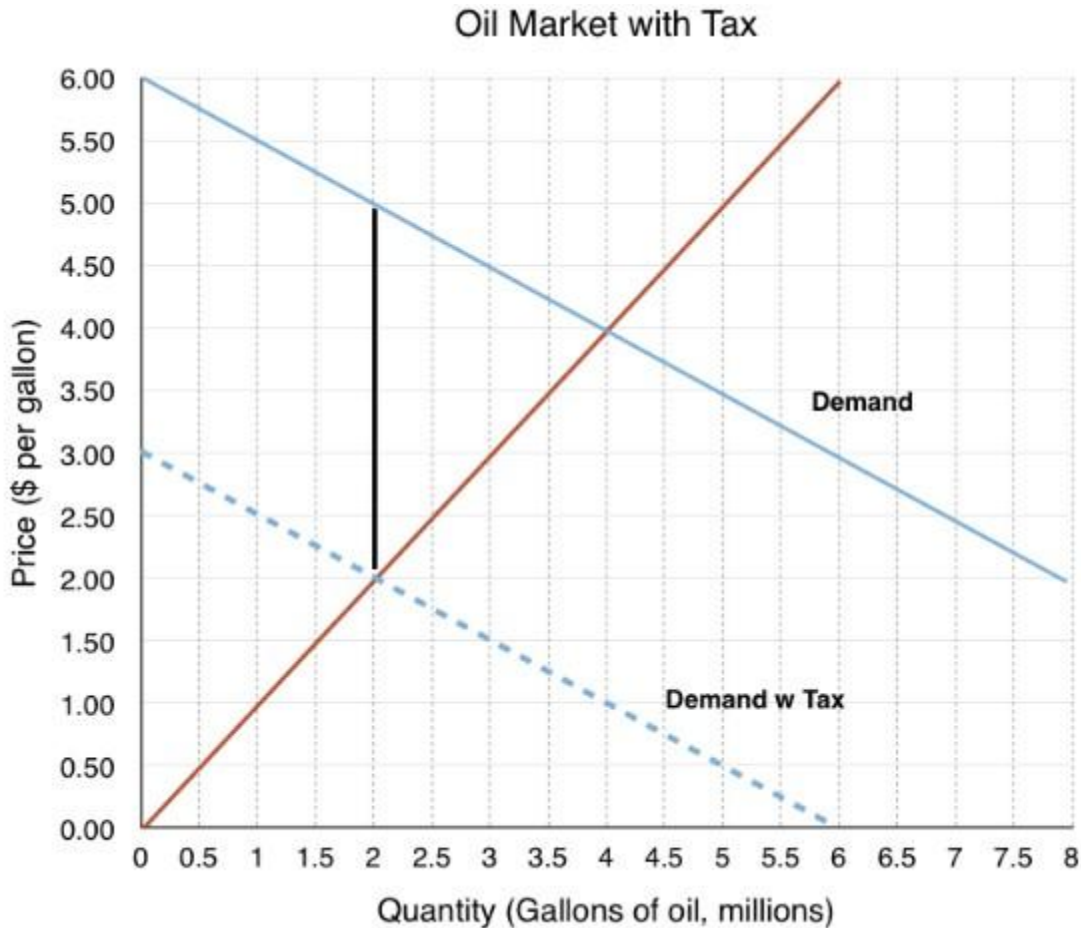


Figure 2: Oil Market with Tax on Consumers

### 4.3.3 Surplus Value

Deadweight loss is a difference in market surplus, so calculating it requires calculating market surplus. Market surplus is the social value generated by equilibria prices, since producers and consumers transact at a price that differs from their reservation prices.

If there be an exchange, there are two alienations, each of which has its separate advantages. This advantage for each of the contracting parties is the difference between the value which they put upon what they give up, and the value of what they acquire. In each transaction of this kind, there are two new masses of enjoyment. In this consists the advantage of commerce (Bentham, 1864, II).

Equilibria models offer a means of calculating this “advantage of commerce.” In our pretax oil market, several million gallons of oil would have been purchased at prices higher than the market

price of \$4. Similarly, producers would have supplied gallons at prices lower than \$4. In that market pretax, we can draw a horizontal line through the equilibrium point, creating two triangles. The area of the top triangle captures consumer surplus, and the area of the bottom triangle captures producer surplus. Together, these triangles represent all the surplus value consumers and producers receive by trading at the market price, which is different from their reservation prices.

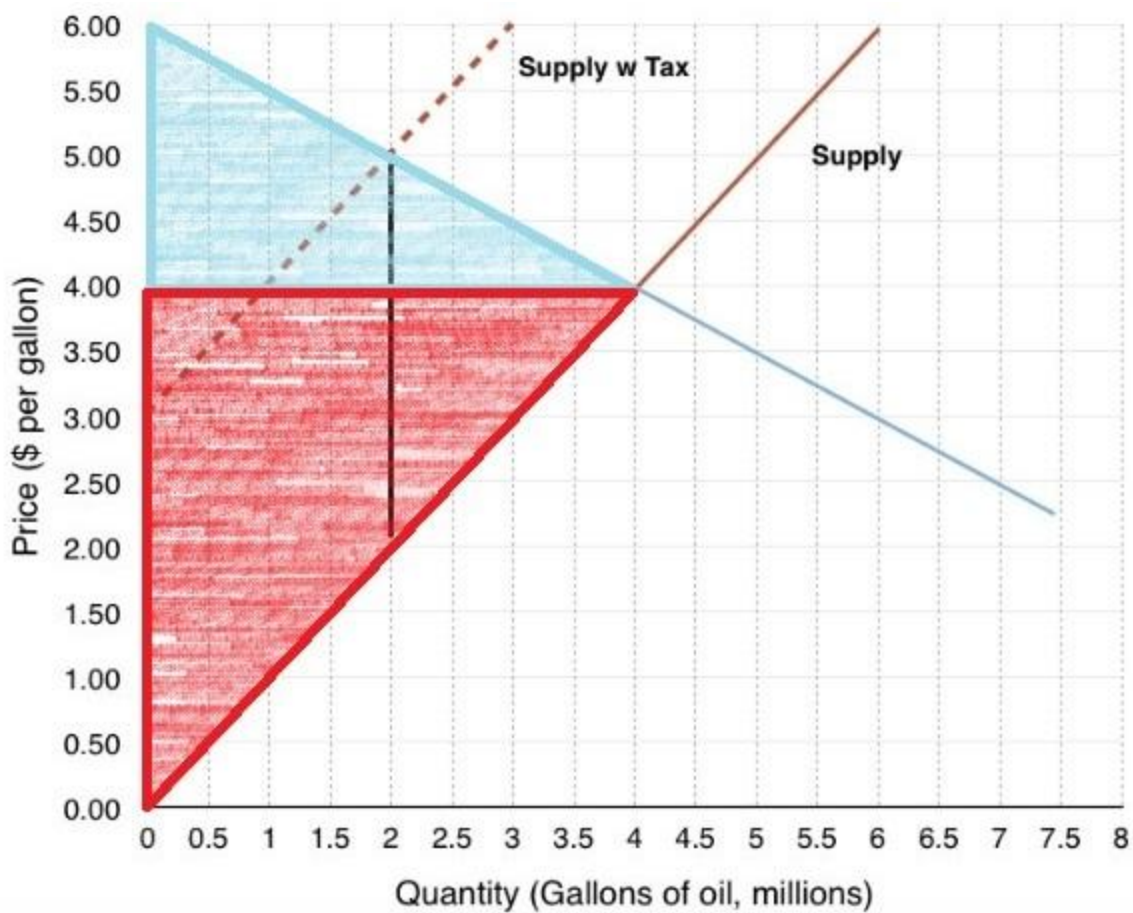


Figure 3: Market Surplus Without Tax

The value of this market surplus can be calculated, since it is simply the sum of the area of the two triangles. The formula for the area of a triangle is  $1/2 \times \text{Base} \times \text{Height}$ .<sup>52</sup> Consumer surplus is  $\frac{1}{2} \times \$2 \times 4,000,000 = \$4,000,000$ . Similarly, producer surplus is  $\frac{1}{2} \times \$4 \times 4,000,000 = \$8,000,000$ . In other words, at the market price, consumers were willing to pay more and producers were willing to supply for less, for a total benefit of \$12,000,000, with two thirds captured by producers.

#### 4.3.4 The Tax Wedge

If taxes merely took some portion of that \$12,000,000 and redirected it to government, that would not constitute a loss in efficiency—the loss of value would not be “deadweight.” But the concern with deadweight loss is that it constitutes a net reduction, not a redistribution, of the social benefit of markets. In the taxed oil market, there are consumers willing to buy at \$2.50, \$3.50, \$4.50 and so on, and producers who want to sell at those prices. Those transactions would create mutual benefit but, because of the tax, they simply will not occur. The value of these lost transactions can be quantified by calculating the area of the wedge the tax creates between the supply and demand curves.

---

<sup>52</sup> More clearly, consumer surplus is  $\frac{1}{2} \times \Delta P \times Qd$  where  $\Delta P$  is the “base” or difference between the highest buying price and the market price and  $Qd$  is the “height” or quantity demanded at equilibrium. Producer surplus is the same, except the “base” ( $\Delta P$ ) is the difference between the market price and the *lowest* selling price.

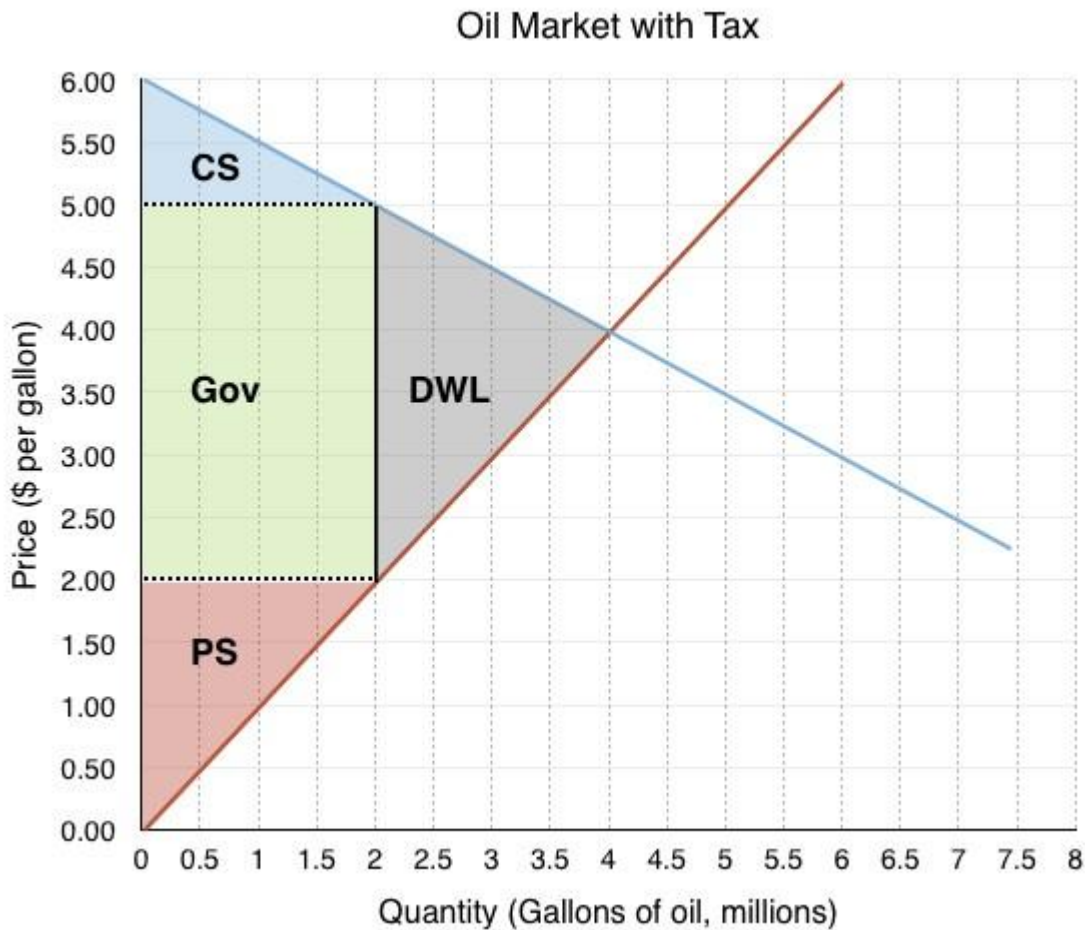


Figure 4: Modeling Tax as Deadweight Loss

Compared with the free market, we can see that the \$3 per gallon tax now divides the market surplus into four parts rather than two. The consumer surplus is now  $\frac{1}{2} \times \$1 \times 2,000,000 = \$1,000,000$ . The producer surplus is now  $\frac{1}{2} \times \$2 \times 2,000,000 = \$2,000,000$ . Government revenue is  $\$3 \times 2,000,000 = \$6,000,000$ . Note that the total surplus among consumers, producers, and government is \$9,000,000. The free-market surplus was \$12,000,000 and so \$3,000,000 in value has been lost due to the tax. This can also be calculated directly as the area

of the deadweight loss triangle<sup>53</sup>:  $\frac{1}{2} \times \$3 \times 2,000,000 = \$3,000,000$ . This value is simply erased by the tax, as it reduces the size of the market in oil.

#### 4.3.5 The “Market” Baseline

Models like the market in oil imply that less tax means more beneficial transactions. In the best case, no tax at all results in the maximum market surplus. As described above, this thinking is generalized to all areas of taxation. For example, it is evident in the orthodox view that increasing the income tax rate inevitably causes deadweight loss. All such views assume a market without tax, and that market provides the baseline for comparison with the reduced surplus achieved in the post-tax market. However, it is unclear how a “pretax” or free market condition provides a meaningful baseline outside of a model context. I want to suggest that it rests on a similar conceptual confusion to that which undermines the notion of full entitlement to pretax income.

What could be meant by a pretax market in oil? I earlier referenced Fleurbaey’s claim that “laissez-faire is in the set of feasible allocations (“no tax” is a particular tax policy)” (2018, p. 42). The idea of a pretax market does seem to rely on something like that claim. Indeed, it is common to think of the state as a kind of external overseer or regulator of a capitalist economy. The state “surveys the outcome of market processes and finds them wanting” (Ackerman, 1977, p. 1). In response, it tries to improve things “taxing here, subsidizing there, regulating everywhere” (Ackerman, 1977, p. 1). In other words, the state reviews the economic outcomes generated by markets and then intercedes on behalf of the public. As Kymlicka put it, “the

---

<sup>53</sup> This is an example of a “Harberger’s triangle” named after Arnold Harberger (1971).

welfare state is primarily concerned with the *post factum* correction of market inequalities, through tax and transfer schemes” (2001, p. 89).

However, this way of conceiving of the state and markets seems incoherent if markets and their economic outcomes depend on institutional features that rely on taxation. Economists, as Demsetz observed, usually take the bundle of property rights as a datum” (1967, p. 347). But legal property rights (along with other legal rights) are essential to the operation of modern markets. Even free market enthusiasts like Posner admitted that “it is very difficult to think about the market apart from the property rights, tort-liability rules, and contractual-enforcement provisions by which the common law regulates the market. In particular, if there were no property rights, and no tort or criminal laws to deter invasions of them, it would be difficult to see how the market would work at all” (1981, p. 790). Today, contemporary thinkers continue to analyze the various ways in which “market economies revolve around contracts and property rights” (Pistor, 2019, p. 47).

Given the fundamental role of malleable legal rights in market operations, the law does more than merely regulate the market. Instead, the state plays a constitutive role in markets as it defines and legitimates the pieces, moves, and powers in an economy. These rules constitute what a market is, rather than merely regulate the operations of some external process. As chess without rules is not “free chess” but rather no game at all (Hollis & Smith, 1994, p. 244), so too is the concept of a modern market free of state intervention a conceptual confusion. Unlike chess, however, the rules of the market are subject to ongoing interpretation, reinterpretation, and revision. In this way they are more similar to the rules of grammar, without which there is not a “free language,” but again, no language at all (Hollis, 152-153). So conceived, markets are

dynamic institutional entities, and so inevitably matters of human creation and political expression (Polanyi, 1944, p. 139).

“Capitalism (as Marx and the classical economists recognized) is not simply an economic system but also a social and political system. For essential to capitalism is the political specification and legal enforcement of a complicated system of extensive property and contract rights and duties that are all conditions of economic production, transfer, and consumption, as well as the political specification and enforcement of laws and regulations needed to maintain the fluidity of markets and competitive enterprise (laws restricting monopolies, price fixing, and other restraints on trade)” (Freeman, 2011, p. 22).

State control over the legal structure of markets largely determines the character and value of transactions. In the example case, the value of a gallon of oil is derived from a vast, complicated institutional structure, regulating everything from the leasing of rights on federal lands to the exploration, production, processing, importation, exportation, and transportation of related goods. Moreover, the entities engaging in these activities are legally defined corporations or other legal persons, each bringing an endless network of rights and duties elaborated in corporate, contract, and tort law (and much more). Changes to these rules would massively influence the value of the available rights related to oil.

Given the central role of the state in constructing markets, we cannot coherently imagine a market in oil that is just like an existing market but without tax. Deadweight loss analysis asks us to do just that, and indeed to afford significant normative weight to a comparative baseline in which the market operates without taxation. It offers a means of estimating the market surplus lost when comparing the no-tax hypothetical with the operations of a real market. Yet if we take the notion of a no-tax baseline seriously, it makes little sense to imagine a world in which everything stays the same except taxes disappear. In fact, without taxes the market would cease to exist in any recognizable way. Perhaps the state would devolve into various tribal groups warring over territory and other resources. Compared to that state of affairs, the market surplus

generated by a market reliant on taxation certainly suffers no deadweight loss.<sup>54</sup> Indeed, it offers far more opportunity for mutually beneficial exchanges than the comparative no-tax baseline. But the crucial point is not what exactly a no-tax economy would look like. It is only that it would not look exactly like the economy we have, minus the taxes. Note that I do not mean to suggest these complexities of real markets have been overlooked. Of course, economists are well aware that real markets are far more complex than models in introductory textbooks. My criticism is rather that any baseline distribution used to calculate deadweight loss is incoherent when it is on the one hand characterized as “pretax” and on the other hand fundamentally relies on institutional features derived from the state, and therefore taxation. The oversimplified example of a market in oil makes this point in a way that even a sophisticated representation of some market cannot avoid.

#### **4.3.6 Deadweight Loss Conclusion**

It is generally assumed that taxes always result in some efficiency loss. The reasoning is straightforward: without taxes, there would be more mutually beneficial economic activity. With taxes, that activity does not occur, and that value is lost. Yet it makes little sense to compare the economy “post-tax” with a hypothetical “pretax” market. Taxes are part of the institutional framework that defines the rules and thereby constructs the market. A measurement of “loss” between a real market and a fundamentally misleading hypothetical cannot provide a consideration of serious normative significance when evaluating tax policy.

---

<sup>54</sup> Epstein makes the following case, “Starting from the original set of entitlements, the formation of the state looks very much like a Pareto-superior move in which some natural rights are surrendered in exchange for the greater protection of the natural rights that are retained. Taxation is the form of surrender, and security is the good received in exchange” (cited in Sampford & Wood, 1991). He is right that the modern state is an improvement compared with the no-state alternative. But the idea that taxation reflects the natural rights surrendered is peculiar, since there would be no natural rights to the economic products of the market-state prior to that state.

#### 4.4 Chapter 4 Conclusion

This chapter considered two economic concepts central to the evaluation of efficiency in taxation: the notion of a Pareto improvement and that of deadweight loss. I argued that in practice both concepts risk undetected and unjustified reliance on dubious baseline assumptions. In the first section, I illustrated how the Pareto principle has been mistakenly thought to provide reason to prefer Pareto-improvements over the status quo even over Pareto-incomparable alternatives. But I also argued that the Pareto-principle provides no basis for this preference, since all it recommends is preferring Pareto-superior alternatives over Pareto-inferior ones. Privileging the status quo is a separate normative commitment, yet one which goes undefended as it is camouflaged in the largely uncontroversial normative significance of the Pareto principle. In the second section, I argued that anxiety about deadweight loss may be misplaced, if it rests on comparing real markets with conceptually confused hypotheticals. Deadweight loss implies that without taxation the quantity of transactions in a market would increase. In fact, taxation is essential in the state construction of modern markets. Without taxes, and therefore without the state, the market would look nothing like a modern market but without taxes. However it would look, it is not clear why that hypothetical would provide a baseline of normative significance. It does not represent a real loss of value, since an existing market “minus taxes” would not be the reality of a no tax economy.

## Chapter 5: The Baseline Problem in Practice

### 5.1 Chapter 5 Introduction

Tax analysis in Canada, as everywhere, is dominated by economists engaged in quantitative research using complex formal models. Naturally that dominance is not absolute, and the study of taxation remains intensely interdisciplinary. There is an inexhaustible supply of analysis from historians, legal theorists, political scientists, sociologists, and more. Yet in the world of policy and administration, the eminence of economic thinking about taxation is undeniable, and moreover it is just one instance of a broader trend. A former Yale Law Dean summarizes the status quo in the following way:

One might without exaggeration say that economics is today the science of administration and policy-making – that every weighing of costs and benefits in pursuit of the greatest good (which is what policy-makers are expected to do) is either a form of economic analysis or an ad hoc judgment with no discipline at all (Kronman, 2017, p. 222-223).

Despite its technical details, there remains widespread recognition that “[t]axation is not a technical matter. It is preeminently a political and philosophical issue, perhaps the most important of all political issues” (Piketty, 2014, p. 630). And so it is that even the most influential economists agree that taxation must be fair, while also admitting that fairness “is not in the end a question of economics” (Slemrod & Bakija, 2017, p. 60).

The question of fairness may be alien to the methods of economics, but it is the animating force of a great deal of political philosophy. What does this imply for philosophers? It seems to imply at least an important interface between taxation and their philosophical efforts. Indeed, it may impose a binding summons on theorists of distributive justice since “tax policy is where we negotiate the relationship between wealth and poverty, equality and inequality” (Heaman & Tough, 2020). In considering these questions, some philosophers have cautioned against

evaluating taxation from too high a level of abstraction. After all, tax systems—and Canada’s is no exception—are invariably technical and arcane. A printed copy of the *Income Tax Act* alone runs over two thousand pages (for example, Thorsteinssons, 2019). It is against this background that Martin O’Neill and Shepley Orr introduce their recent collection of philosophical perspectives on taxation with the following admonition:

Thinking clearly about tax is not an easy task, as much that is of central importance is missed if one proceeds at too great a level of abstraction, and issues of conceptual and normative importance often only come sharply into focus when viewed against real-world questions of implementation and feasibility (2018, p. 2).

They are not alone in emphasizing the real, the concrete, and the feasible in thinking about taxation. Alan Hamlin substantiates their concern, observing that political philosophy pays insufficient attention to the details of taxation:

Taxation is often regarded as a ‘black box’ technology that can be called upon to put into effect whatever distribution of economic benefits and burdens that is required by the normative theory under discussion, without the need to consider the more detailed internal properties of the black box itself (2018, p. 17).

He goes on to argue that taxation may be non-ideal by its very nature. Assessing the proper role for taxation in an idealized world of full compliance, perfect knowledge, and unlimited political power may be fruitless. In such a world, there may be no place for taxation at all. Certainly taxation in that world would, at minimum, look very different from our own. In a sense, then, taxation is inherently a theory of the second best. All of this makes a compelling case that philosophers should continue to think about taxation but, at the same time, remain sensitive to a wide array of less-than-ideal concrete details.

However, in this project I have argued that dwelling in the details risks obscuring a fundamental insight, and one that may after all be most visible from a high level of abstraction. That insight is simply that tax justice is “not a matter of applying some equitable-seeming

function to a morally arbitrary initial distribution” (Murphy & Nagel, 2002, p. 30). I have described how the standard practice of tax policy analysis is in direct conflict with this insight. It takes the legal distribution of taxable income (for example) as a baseline to which an “equitable-seeming function” is applied (see, for example, Duff, 2018). In consequence, resulting evaluations of fairness are meaningless unless individuals are morally entitled to the legal distribution of taxable income. But in the standard approach to assessing tax fairness, that indispensable consideration is simply swept aside. With some assumed baseline distribution in place, tax fairness is treated as an autonomous domain of moral assessment. The significance of that domain for justice more broadly is nearly always recognized: “there can be no fairness without tax fairness” (Heaman & Tough, 2018). But the less often appreciated point is that there can be no tax fairness without background fairness. Again, I point to Arthur Pigou’s (1947) mostly ignored insight that,

To hold that the law about taxes ought to affect different people’s satisfactions equally, while allowing that the rest of the legal system may properly affect them very unequally, seems not a little arbitrary.

I think the deepest threat to justice in taxation, then, is what we take for granted, what we accept as a distributive baseline. It is, as I have argued, a baseline problem.

From this vantage point we can see how taxation poses a unique peril to distributive justice: it is the state’s most powerful tool for establishing ‘right by fact.’ I refer here to Jean-Jacques Rousseau’s memorable complaint of Hugo Grotius that “[h]is usual method of reasoning is constantly to establish right by fact” (1762). The Grotian tendency involves an excessive deference to authority, and in particular to the status quo. Rousseau argues that no method “could be more favorable to tyrants,” since the Grotian method conceals the injustice of fact by transforming it with a veneer of right. The standard analysis of taxation is complicit in a modern

version of just such a Grotian transformation. By taking for granted the moral status of some initial distribution, the standard approach to tax fairness does indeed conceal the arbitrariness of that initial distribution. But it also preforms a full metamorphosis of fact into right, since taxation necessarily involves the transformation of a factual pattern of holdings into a set of juridical entitlements which, in the minds of most, have significant legitimacy. The process obscures the dubious normative credentials of some set of initial holdings, but it also applies an equitable-seeming formula to that distribution, thereby ostensibly transforming fact into right.

All of these points have been illustrated with various theoretical examples but there is also, as O'Neill and Orr emphasize, the practical reality of taxation. To illustrate how these ideas connect directly with this reality, I now consider two case studies. The first is the now entrenched public practice of reporting tax expenditures in both the U.S. and Canada (a practice that also reaches around the globe). The second is the Canadian public discourse about tax justice, and in particular the consensus view that 'paying a fair share' is at its heart. Together I hope these two cases help illuminate how public policy about taxation is infected with a baseline problem.

## **5.2 Tax Expenditures**

### **5.2.1 Introduction**

Since its inception, tax expenditure analysis has been mired in controversy. The most prominent criticism is that it relies on normative judgments and so violates the constraints of neutrality and objectivity imposed by public policy. There are two important responses to this criticism. The first argues that the institutionalized best practice of tax expenditure analysis is more neutral than critics allow. The second accepts the criticism but proposes a new paradigm for tax expenditure analysis explicitly designed to satisfy constraints of neutrality and objectivity.

In this section I argue that neither response succeeds, recommending candid normativity as an alternative. I begin by introducing the concept of a tax expenditure and illustrating its significance in public finance. Two key points are that (1) the intuitively simple idea is more enigmatic than it first appears and that (2) it essentially relies on a normative baseline. These points are then applied to the Department of Finance Canada's analytic framework, exposing numerous normative judgments that characterize its Benchmark Tax System. The defining attempt to avoid such judgments is the 2008 Joint Committee on Taxation's proposal for reform. I therefore introduce and assess that proposal, concluding that it is merely cryptonormative rather than genuinely neutral and objective. In response I suggest rejecting the conventional anxiety about 'subjective' normative judgments, accepting instead that tax expenditure analysis is a thoroughly normative project.

### 5.2.2 Spending by Not Taxing

A clergy member in British Columbia with a salary of \$70,000 will be taxed at a marginal rate of around 30%.<sup>55</sup> If the government wishes to increase the clergy member's after-tax income by \$3,000 it can do so in two ways. The first is to issue him a taxable payment of about \$4,300, which it can do by depositing \$4,300 directly in his account or by cutting him a check for the same. The second way to increase his spendable income by \$3,000 is to reduce his taxable income by \$10,000. In the first case resources are transferred from the clergy member to the state via taxation and then transferred back to him via direct payment. In the second case a tax deduction simply circumvents that back-and-forth process. Of note is that the economic consequence for both him and the government is identical both cases. In fact, the Government of Canada has opted for the latter approach, permitting clergy to deduct up to \$10,000 in living expenses from their income, pursuant to 8(1)(c) of the *Income Tax Act*.

The clergy residence deduction is an example of a tax expenditure. The term 'tax expenditure' was introduced to the public by the eminent American tax scholar Stanley Surrey in a 1967 speech while he was Assistant Secretary of the Treasury for Tax Policy (Hellmuth & Oldman, 1973). He promoted the concept for the remainder of his life: "A dollar is a dollar—both for the person who receives it and the government that pays it, whether the dollar comes with a tax credit label or a direct expenditure label" (1970, p. 717). His final work titled *Tax Expenditures* was published only months after his death in 1985 (McDaniel & Surrey). The central thesis of his project was that in many cases the economic effects of direct government spending are indistinguishable from favourable tax treatment, and so 'spending' via tax preference is a central function of existing tax regimes:

---

<sup>55</sup> See (Wolfman, 1985) for an American analogue.

The federal income tax system consists really of two parts: one part comprises the structural provisions necessary to implement the income tax on individual and corporate net income; the second part comprises a system of tax expenditures under which Governmental financial assistance programs are carried out through special tax provisions rather than through direct Government expenditures. This second system is grafted on to the structure of the income tax proper; it has no basic relation to that structure and is not necessary to its operation. Instead, the system of tax expenditures provides a vast subsidy apparatus that uses the mechanics of the income tax as the method of paying the subsidies (1973, p. 6).<sup>56</sup>

In short, if a government wants to increase the resources of some targeted group, it can either spend directly on that group or it can reduce their taxes owed (OAG BC, 2018, p. 7).

Tax expenditures are a significant feature of public finance. Some, such as the clergy residence deduction, are small and targeted. The Department of Finance Canada (DFC) estimates that in 2022 that deduction will only be claimed by about 25,000 individuals and will cost \$105 million in lost personal income tax revenue, a relatively minor sum on a national balance sheet (DFC, 2021). But the Canadian government chooses to increase spendable income by reducing taxes owed in a number of cases which have a more significant budgetary impact. For example, the DFC estimates 3.5 million families received the Canada Child Benefit in 2020, at a cost of over \$27 billion in lost revenue. Another significant measure is the partial inclusion of capital gains, from which about 2.9 million individuals, 254,000 corporations, and 37,000 trusts reportedly gained about \$21 billion in 2018. The tax deferral option provided by the Registered Retirement Savings Plan is estimated to benefit 8.9 million households annually to the tune of

---

<sup>56</sup> His final work in 1985 (p. 3) put it this way: “The tax expenditure concept posits that an income tax is composed of two distinct elements. The first element consists of structural provisions necessary to implement a normal income tax, such as the definition of net income, the specification of accounting rules, the determination of the entities subject to tax, the determination of the rate schedule and exemption levels, and the application of the tax to international transactions. The second element consists of the special preferences found in every income tax. These provisions, often called tax incentives or tax subsidies, are departures from the normal tax structure and are designed to favor a particular industry, activity, or class or persons. They take many forms, such as permanent exclusions from income, deductions, deferrals of tax liabilities, credits against tax, or special rates. Whatever their form, these departures from the normative tax structure represent government spending for favored activities or groups, effected through the tax system rather than through direct grants, loans, or other forms of government assistance.”

about \$16 billion. Total tax expenditures enumerated in Canada’s federal report exceed \$100 billion, and in addition the provinces also report expenditures of several billion each (Lester, 2019, p. 756).<sup>57</sup> For comparison, Canada’s entire federal expenditure for 2019 (pre-covid) was \$355.6 billion.<sup>58</sup>

### 5.2.3 Spending in the Shadows

Why care about tax expenditures? They are thought significant primarily because they are one of the more shadowy mechanisms through which the state allocates resources. The core insight guiding their analysis is that “certain provisions of the tax laws are not really tax provisions, but are actually government spending programs disguised in tax language” (Thuronyi, 1988, p. 1155). The reform program of removing that disguise continues to exert a powerful influence on public finance. For example, in 2018 the Auditor General of BC affirmed that “[i]t is important to consider tax expenditures and direct spending together to ensure that the total costs of government programs and services are understood” (OAG BC, p. 7).

Her point about ensuring an understanding of “total costs” recalls the same considerations that animated Surrey in the 1970s. In addition, he emphasized that, despite their economic identity, tax expenditures differ from direct spending because they receive less public scrutiny and are rarely subject to cost-benefit analysis. The Auditor General evidently agreed, expressing concern that “[t]ax expenditures are debated and voted on only when the legislation for them is enacted or changed. This doesn’t happen very often.” She highlighted some current expenditures which are very old, such as the homeowner grant introduced in 1957. Economic and social

---

<sup>57</sup> Reports from the United States standardly enumerate well over 1\$ trillion in tax expenditures and “have often exceeded direct, discretionary federal spending (Fleming & Peroni, 2008, p. 439).

<sup>58</sup> <https://www.budget.gc.ca/2019/docs/plan/budget-2019-en.pdf>

conditions have changed significantly since the 1950s, and the public today may have a quite different view about the need for various tax measures. In general, she restated Surrey's earlier concern that direct spending receives significant analysis and public scrutiny in the budgeting process, whereas tax expenditures are generally ignored. The effects of this opacity are not difficult to imagine. Consider again the clergy residence deduction. The stated policy objective is to "recognize the special nature of the contributions of members of the clergy." A budget line item recognizing these "special contributions" by direct depositing \$105 million in the bank accounts of clergy members might raise the public eyebrow during the budgeting process, whereas a 1950s-era deduction of clergy rental payments may tend to go relatively unnoticed. But if tax preferences are just a more covert means of welfare spending, then this difference in response may be attributed to a lack of transparency rather than any difference in principle.

The central concern motivating calls for reporting and reform is therefore transparency, but there are others as well (Fleming & Peroni, 2008). There are equity concerns, since tax expenditures are often deductions and so worth more to those with higher incomes (so-called 'upside down' subsidies). Others are concerned about ease of administration since exceptions to general rules tend to make taxation more complex. Still others think tax expenditures are less predictable or more subject to political inertia than other forms of spending.<sup>59</sup> A recent study by David Pozen (2007) summarizes several of these concerns:

Tax expenditures increase tax complexity and compliance burdens (administrative concerns); they are more opaque and esoteric than direct expenditures and so relatively more immunized from public debate or scrutiny (visibility concerns); they are more likely than direct expenditures to have regressive first-order effects (equity

---

<sup>59</sup> Some of these concerns are older than the term 'tax expenditure' itself. Blum (1955, p. 251) writes: "if it be decided to subsidize a certain activity, we should be hesitant about administering the subsidy by way of a tax preference. Subsidies in this form vary directly in amount with the tax brackets of the recipients; they are invariably hidden in technicalities of the tax law; they do not show up in the budget; their cost frequently is difficult to calculate; and their accomplishments are even more difficult to assess. Partly for these very reasons they are likely to become fixtures which are not easily removed."

concerns); their uses may evolve in unexpected ways (predictability concerns); and they may stimulate unproductive activity or reward behavior that would have occurred anyway (efficiency concerns).

In general, critics are not necessarily against the economic effects of tax expenditures: the Canada Child Benefit might be a wholly defensible public policy. But, they say, that should be decided in the full light of public inquiry, rather than in the shadows of tax legislation.

Tax policy analysis was immediately gripped by the idea of tax expenditures and reporting them is now a widespread practice. Indeed, though the practice of reporting tax expenditures began in the United States, it is now a global phenomenon. The United States Treasury first published a list of tax expenditures in 1967, and by 1974 it prepared special analyses of its use of the tax system to achieve policy goals. These reports were mandated first by the *Congressional Budget Act* and soon after by the *Tax Reform Act*. The Department of Finance Canada followed, first issuing a tax expenditure report in 1979, a now annual practice which it describes as “a key component of the Government’s reporting on the federal tax system.” Today both the International Monetary Fund (IMF) and the Organization for Economic Co-operation and Development (OECD) emphasize tax expenditure reporting as a best practice for public finance accounting.<sup>60</sup> The latter is directly involved in tracking and assessing the tax expenditure reports of member nations.<sup>61</sup> In total, the Global Tax Expenditure Database (G-TED) lists over 100 countries with publicly available data on tax expenditures, summing to trillions of dollars in lost government revenue.<sup>62</sup>

---

<sup>60</sup> International Monetary Fund, *Manual on Fiscal Transparency*, Fiscal Affairs Department, 2007; Organisation for Economic Co-operation and Development, *OECD Best Practices for Budget Transparency*, 2002.

<sup>61</sup> [https://www.oecd-ilibrary.org/governance/tax-expenditures-in-oecd-countries\\_9789264076907-en](https://www.oecd-ilibrary.org/governance/tax-expenditures-in-oecd-countries_9789264076907-en)

<sup>62</sup> <https://gted.net/>

#### 5.2.4 The Tax Expenditure Enigma

Public finance looks quite different through the lens of tax expenditures, but many have noticed that the very concept seems enigmatic. One “unsettling” thought is that the concept implies that a state can spend more by taxing less (Prasad, 2011, p. 254). It is at least odd to hold that a state that taxes not at all is engaged in extensive spending through its wholly restrained regime of (non) taxation. At the limit, a fully general equivalence between spending and non-taxation seems to imply that the largest welfare state is equivalent to the absence of any state at all. Of course no one defends that final claim, but Prasad argues that it follows from the internal logic of tax expenditure analysis. And she is right that the (partially applied) logic of spending by not taxing does animate the tax expenditure literature.

A clear example of this logic comes from Christopher Howard’s famous 1997 study in which he argues that tax expenditures in the United States constitute a “hidden welfare state,” amounting to trillions of dollars in subsidies. He traces the history of four significant tax expenditures: mortgage interest, employer pensions, the Earned Income Tax Credit, and the Targeted Jobs Tax Credit. His history supports concerns about transparency, since it reveals that the political struggle for tax expenditures differs significantly from direct spending. Tax expenditures draw less scrutiny and input from interest groups, and they draw more support from conservative politicians than traditional social programs. But he emphasizes that their core function is to provide benefits without direct spending.

Others, most notably Jacob Hacker (2002), later developed Howard’s thesis, arguing that tax expenditures on health and pensions provide a private welfare state that frustrates the public provision of goods through direct spending. These authors agree that the logic of tax expenditures unveils a government much larger than it appears on the budget. As Howard put it,

“the American welfare state is actually 40 percent larger than commonly believed” (cited in Prasad, 2011, p. 252). And similar claims are often made about other governments. For example, Munir Sheikh (2014, p. 22) concludes that accounting for tax expenditures reveals that the “true size” of the Canadian government is about 25% larger than is commonly recognized. These arguments claim that the welfare state is larger than it appears precisely because it taxes less.

The logic of tax expenditures also has implications for the politics of government spending. Claims about ‘fiscal responsibility’ and ‘tax-and-spend’ are only evaluable in light of a complete public accounting, including both direct and indirect spending. If the traditional “dichotomy between tax and spending programs is false” (Thuronyi, 1988, p. 1156), then a ‘fiscally responsible’ conservative party might turn out to spend more than a ‘tax-and-spend’ liberal party once both platforms are brought under the tax expenditure lens.<sup>63</sup> Indeed, one study found that accounting for tax expenditures revealed that all political parties in power ‘spend’ around the same.<sup>64</sup> Accordingly, some liberals claim that conservatives spend much more than they admit—they simply do it through the shadowy mechanism of targeted (non) taxation rather than in the full light attracted by direct spending. As Howard put it, after accounting for tax preferences the American Republican party “doesn’t exactly sound like the party of small government” (2009, p. 96). Perhaps conservatives can concede the point. After all, if they can ‘spend’ more with a smaller state, so much the better. Not surprisingly, therefore, some have

---

<sup>63</sup> “The back door to government subsidies marked ‘tax relief’ is easier to push open than the front door marked ‘expenditures’ or the side door marked ‘loans, guarantees, and insurance.’ Rather than run the gauntlet of the Budget Bureau and the congressional Appropriations Committees, groups seeking subsidies turn to the tax committees of Congress for Government support [ . . . ] The very groups that use this back door are often among the most insistent advocates of responsible and informed Government budgeting. Yet here is a whole catacombs of Government benefits which are largely hidden from public view, let alone, periodic review. Once imbedded in the tax structure, the preferential provisions are treated as inalienable vested rights, impervious to changes in tax rates, economic policy, and technology” (Heller, 1959, p. 181).

<sup>64</sup> “The essential problem is this: presidents from both political parties and Congress today use tax exclusions, deductions and credits the way my mother used chicken soup, as a cure-all for every economic and social ill the country faces” (Graetz, 2008, p. 1447).

embraced the view that after accounting for tax expenditures “Americans spend more per capita on social welfare than do any of the Nordic countries” (Fishback, 2010). So in a strange inversion of Surrey’s original argument (that tax expenditures obscure favouritism for wealthy and powerful interests) some have argued that tax expenditures uncover a previously underappreciated benevolence at the heart of American governance. By doing less, it turns out, government does more.

So far all of this suggests that Prasad’s concern is well-motivated: if government can increase in size by spending less, then we are left with the dubious conclusion that the largest state may be no state at all. What has gone wrong in the analysis? The problem, she thinks, is uncritical acceptance of the thesis that direct spending is equivalent to tax preference. At first, this equivalence thesis seems undeniable: our previously discussed clergy member and the government are both in the same economic situation whether the member’s spending power is increased by a direct deposit (“spending”) or a deduction on his taxes (“tax preference”). This and other compelling observations display the “simple, intuitively obvious idea” at the “core” of all tax expenditure analysis (Fleming & Peroni, 2010, p. 136). Yet examples like these rely on a more general principle, namely that *giving* (direct spending) is the same as *not taking* (tax preference). Of course these are not *always* the same, since if I give you a car I have done something quite different from not taking your car. But are they sometimes the same?

Prasad observes that there may be only one condition under which giving and not taking are equivalent, and that is when the giver has a right to the resource at stake. If I have a right to a car in your possession, and I want you to have it, my giving it to you is equivalent to foregoing my right to take it back. In this case, giving and not taking are the same. If you have a right to the car, however, I am obligated not to take it and have no right to give it. In this case, giving and

not taking are different. Any equivalence between giving and not taking therefore depends on the distribution of entitlements. Thus the veracity of the equivalence thesis in any tax expenditure case depends on a right of the state to take the resources at stake. If the equivalence thesis is always true, or true in general, that can only mean that the state has a right to all resources.

### **5.2.5 The ‘Last Penny’**

If tax expenditure analysis depends on the implausible claim that the state has a right to all resources, then that constitutes a significant criticism of the project. Indeed, the American Joint Committee on Taxation (JCT) reviewed the literature on tax expenditures and found that some critics were skeptical of the analysis believing it rested on just that claim.<sup>65</sup> It called this criticism the ‘Last Penny’ argument, a phrase taken from Charles Fried’s allegation that tax expenditure analysis accepts the “sinister premise .... [that one should] think of all income as virtual state property, and forbearance to tax away every last penny of it as itself a tax expenditure” (JCT, 2008, p. 35). Similarly, a 1999 report for the Joint Economic Committee of the U.S. Congress opens by framing tax expenditures as assuming that: “provisions that shield components of [a] broad definition of income are viewed as depriving the government of its rightful revenues; these lost revenues are regarded as properly belonging to the federal government” (Hall, 1999, p. 1). Naturally, concerns of this kind are older than the term ‘tax expenditure’ itself.<sup>66</sup> For example, Neil Brooks relates an account of a disagreement among British Parliamentarians in 1863. Tory member William Gladstone argued that charitable contributions should not be exempt from the income tax. After all, he said, if the state wants to

---

<sup>65</sup> The U.S. General Accounting Office made a similar finding: “Some observers believe that labeling these provisions tax ‘expenditures’ implies that all forms of income inherently belong to the government” (1994, p. 2).

<sup>66</sup> One history of the concept underlying the term ‘tax expenditure’ is (Forman, 1986).

spend revenue on charities why not maintain control over that expenditure? In reply, Stafford Northcote pilloried Gladstone's framing of a deduction as a state expense: "The right hon. Gentleman, if he took £5 out of the pocket of a man with £100, put the case as if he gave the man £95. . ." (Brooks, 1986, p. 684).

Defenders of tax expenditure analysis are quick to distance themselves from the 'last penny' commitment. Burman calls it the "most peculiar argument" (2003, p. 620), the JCT finds it "wholly without merit" (2008, p. 36), and it is "dismissed" by Clifton and Peroni (2010, p. 175) as the "government-owns-all-your-income canard" (2008, p. 496). The essence of their response is that the 'last penny' argument is only credible if tax expenditure analysis "can be fairly understood as asserting that Congress has a normative obligation to adopt a generally applicable income tax and that 100% is the normatively correct rate for such tax" (Clifton & Peroni, 2008, p. 492). On the contrary, however, tax expenditure analysis is not inherently committed to any particular description of the income tax. As Victor Thuronyi explains:

[The 'last penny'] criticism, of course, misses the point. Tax expenditure analysis does not assert that the federal government has a right to tax away everyone's income. Rather, it asserts that the government taxes to the extent of the normative tax, but gives back a substantial amount of this tax revenue through various tax based spending programs (1988, p. 1178).

Prasad argued that giving and not taking are only equivalent when the giver (or not taker) has a right to the resource at stake. This appears to imply that tax expenditure analysis is subject to the 'last penny' criticism. But if tax expenditure enthusiasts can explain why the state has a right to some, but not all, of the resources at stake in taxation, then they can agree with Prasad that giving and not taking are not always equivalent. Whether Stafford Northcote is right about the non-equivalence of taking £5 of £100 and giving £95 depends on the particular case and its relation to the entitlement baseline.

Tax expenditure analysis must maintain that not taxing is the same as direct spending in a large number of cases. Again not *all* cases, as it cannot accept a general equivalence between spending and not taxing while at the same time rejecting the ‘last penny’ criticism. What it requires, of course, is a baseline description of taxation to use as a reference point from which to measure the economic effects of various tax measures. This baseline is what Thuronyi referred to in his earlier cited claim that “the government taxes to the extent of the *normative tax*” (1988, p. 1187, emphasis added). Given some normative description of a tax system, it may be possible to assess differences between that baseline and the real tax system without accepting a total state entitlement to all revenue. The success of the analysis therefore turns on the legitimacy of that baseline.

### **5.2.6 The Normal Tax**

By far the most powerful criticism of tax expenditure analysis is that it lacks any defensible baseline description of taxation against which to measure expenditures (Fleming & Peroni, 2010, p. 142). Proposed baselines are described variously in the literature as the normal tax, normative tax, reference tax, benchmark tax, and so on, but they all respond to the same challenge: “[t]he unavoidable problem is that, by definition, tax expenditure analysis requires comparing actual rules to some hypothetical” (JCT, 2008, p. 16). Even at its inception, Surrey accepted that characterizing a “normative tax” was unavoidable: “tax expenditure analysis is based on the concept of a normative tax of the type under consideration” (1985, p. 3). I noted above that Surrey saw the tax system as comprising two components, which he called “structural” and “special” provisions. His aspiration was that analysts could characterize the

normative tax in terms of the structural elements of the tax regime and then measure the economic impact of special provisions as deviations from that baseline.

Critics immediately pointed out that the distinction between structural and special provisions was not rigorously formulated. Surrey's characterization of the normative tax looked less like a neutral description likely to command wide support and more like an articulation of his idiosyncratic preferences. Indeed, the general implausibility of a neutral description of the normative tax baseline was noted even before Surrey introduced the term 'tax expenditure' to the public. One of the earliest critics was Boris Bittker, who declared that Surrey had merely adapted the Haig-Simons definition of income to suit his personal judgments about what to include as structural elements in the normative tax baseline. Any tax on income requires a definition of income in order to identify the tax base, and Surrey's normative tax included a tax on both personal and corporate income. Predictably, then, he would turn to the conventional standard for defining income in tax policy then and still today (Slemrod & Bakija, 2017), which comes from the work of Robert Haig and Henry Simons.

Haig (1921, p. 27): income is "the money-value of the net accretion to one's economic position between two points of time."

Simons (1938, p. 50): income is "the algebraic sum of (1) the market value of rights exercised in consumption and (2) the change in the value of the store of property rights between the beginning and the end of the period in question."

The combined result is the Haig-Simons definition of income, a comprehensive definition that includes all consumption and changes in wealth.

Instead of defining income in theoretically principled Haig-Simons terms, Surrey followed the standard tax administration practice of altering the definition to exclude some consumption and changes in wealth from the tax base. Just as real tax systems commonly exclude inheritance, gifts, and imputed income on owner-occupied dwellings, so too did Surrey's

hypothetical characterization make idiosyncratic exclusions. Bittker was rightly critical of any ambitions to neutrality, noting that Surrey's description of the normative tax brought "some issues to the fore only to conceal others" (1969, p. 244). Later critics have largely agreed with that assessment: "Surrey's normative tax reflected his vision of what the U.S. income tax should look like, given the constraints of politics. By incorporating 'generally accepted' concepts, though, Surrey endeavored to shield this subjective vision behind the authority of a consensus of experts" (Thuronyi, 1988, p. 1166).

Bittker's criticism is probably even more general than he intended. Perhaps he would have been satisfied if Surrey had retracted all adjustments to the comprehensive income tax base. However, the Haig-Simons definition of income is itself an expression of what ought to be relevant in constructing the tax system, namely all consumption and changes in wealth. This is not a 'neutral' characterization of the tax base.<sup>67</sup> After all, many people believe that income taxes are regrettable as a general category, preferring to characterize the normative tax in terms of consumption alone (Hall, 1999). Indeed, Leonard Burman later described this as "the most contentious issue" in defining the normal tax (2003, p. 618). The choice has significant consequences for what counts as an expenditure (see OMB 2009, fn. 17). As Burman notes, the preferential rate on capital gains looks completely different depending on the normal tax baseline. Against an income tax baseline, taxing only half of realized gains is one of the largest national expenditures. However, the capital gains rate would be zero under a consumption tax, and so the partial inclusion of capital gains is a negative tax expenditure: a tax in excess of the baseline rate. So characterizing the normative tax in terms of an income tax at all is a normative choice, and so too are the more specific choices about how income should be defined and taxed.

---

<sup>67</sup> Though some describe it as a "normatively correct income tax base" (Fleming & Peroni, 2008, p. 447).

The JCT review of the literature on tax expenditures came to essentially the same conclusion, finding that attempts to characterize the normal tax are inevitably value-laden and without a rigorous formal basis:

Since the publication of Bittker's [1969] essay, many other observers, representing a wide range of political and policy perspectives, have echoed his criticism of the "normal" tax base as a subjective compilation of judgments without a rigorous formal basis. Despite nearly uniform agreement on this point, however, no consensus has emerged on an alternative approach (2008, p. 28).

The most enduring and fundamental criticism of tax expenditure analysis, therefore, is that it suffers from a baseline problem. At the highest level of analysis there are two popular strategies for rescuing tax expenditures from this problem. The first is to deny that the operative baseline is 'normative' in any concerning sense, and the second is to attempt a neutral reconstruction of the baseline. The first strategy is reflected in the inclusion of tax expenditure analysis as a best practice according to both national and international standards of public accounting. Canada has developed what is considered an exemplary model for reporting tax expenditures, and its analytic framework is illustrative of the first strategy because it reveals an attempt to avert the baseline problem by distancing its approach from 'normative' alternatives. The defining instance of the second strategy is the 2008 JCT's new paradigm for tax expenditure analysis. The JCT conceded that the orthodox analysis relies on a flawed baseline and so attempted to replace it with a non-normative alternative derived from neutral principles of law and economics. The following sections consider both Canada and the JCT's attempts to resolve the baseline problem and explain why neither is successful.

### **5.2.7 The Canadian Framework**

The Canadian federal government is considered a global exemplar in tax expenditure reporting. The DFC sets out its preferred analytic structure along with the results of its analysis

in the annual *Report on Federal Tax Expenditures*. The report begins by introducing what is called the “Benchmark Tax System,” emphasizing the need for a standard benchmark. Instead of an unsystematic practice of reporting (or not reporting) specific tax measures as expenditures, a standard benchmark provides “an objective basis for selecting which tax measures to include in the report” (DFC, 2021, part 1). Of course, this also means that the results of the analysis entirely turn on how that benchmark is defined; a flawed benchmark identifies a flawed list of tax measures and is a paradigm example of a baseline problem.

The DFC report distinguishes between two broad approaches to characterizing the benchmark. The first is the “core function” approach which involves characterizing the benchmark as “only the most fundamental aspects of the tax system.” The alternative is the “normative approach” which involves characterizing the benchmark as is “considered optimal from an economic and tax policy perspective.” Canada has opted for the core function approach because it “ensures that information is reported on a wide range of tax measures” and is “simpler and less subject to interpretation.” As such, the DFC benchmark is based on an understanding of the tax system as having two functions. The first and principal function of the tax system is “to raise the revenues necessary to fund government expenditures.” The second is “to achieve public policy objectives through the application of specific measures such as preferential tax rates, exemptions, deductions, deferrals and tax credits.” These measures are “deviations from the core function of the tax system, at the cost of lower tax revenues.” Thus the report distinguishes between a baseline tax structure necessary to achieve the core function of the tax system, and preferential deviations from that baseline which are then called tax expenditures. So far, this division neatly tracks Surrey’s original description of tax regimes as including “structural” (income generating) and “special” (other policy achieving) tax measures.

### 5.2.7.1 Canada's Flawed Baseline

The DFC adopts Surrey's two-part model of the tax system, but Surrey never managed to deliver a principled distinction between structural and special provisions; his normal tax was always perceived as too idiosyncratic and subjective. Tax expenditure analysis has been caught in a kind of dilemma ever since (JCT, 2008). On the one hand, it is supposed to provide "an effective and neutral analytical tool." On the other hand, "the premise of the analysis (the validity of the 'normal' tax base) is not universally accepted." As a result, the DFC is keenly aware that it must defend its core function baseline against charges of idiosyncratic subjectivity; it must provide some legitimate grounding for its characterization of the benchmark. One might expect, for instance, a principled and normatively defensible distinction between 'core' and 'non-core' provisions.

Instead, the DFC attempts to avoid having its analytic framework characterized as normative at all. Recall that the DFC introduces the benchmark tax system by distinguishing between "core function" and "normative" characterizations of the benchmark, which it describes as follows:

**Core Function:** "only the most fundamental aspects of the tax system"

**Normative:** "[a] tax system that is considered optimal from an economic and tax policy perspective"

But this distinction only obscures rather than justifies the normative judgments guiding the analysis. It is a transparent attempt to distance the preferred 'core function' characterization of the benchmark from 'normative' alternatives and therefore render the chosen benchmark more neutral and objective in comparison. Yet identification of the core functions of a system of taxation is an unavoidably normative project. After all, as I will illustrate below, there are no value-free criteria for determining which aspects of the tax system are the "most fundamental" or

how exactly those aspects should be characterized (DFC, 2021). Contrary to the DFC report, the proposed distinction in approaches is not between the normative and the non-normative, but rather between two normative approaches. The first aims to define the benchmark in terms of the core functions of the tax system, and the second aims to define the benchmark in terms of optimal design. But both are thoroughly normative descriptions of the benchmark, and what is needed is some reason to prefer one over the other.

As noted, the DFC framework includes only two reasons to prefer the core function (or “most fundamental aspects”) approach. The first is that it ensures a lot of tax measures will be reported (breadth). The second is that it is “simpler and less subject to interpretation” than a “normative” alternative (simplicity and precision). This slight justification is offered as a response to the JCT dilemma, which again is that tax expenditure analysis is supposed to be a neutral tool yet is unavoidably grounded in some account of a normative tax. It seems inadequate in response to merely declare the core function approach “simpler and less subject to interpretation” and so implicitly more neutral and objective than a normative alternative. The values it cites (breadth, simplicity, and precision) are plainly relevant for a tool that is intended to be both quantitative and informative. But their value comes from pairing them with a correct baseline. For example, wide reporting on tax measures is of little use if it is uninformative and misleading because the wrong measures are included. Similarly, precision and simplicity are hardly the cardinal virtues of a baseline in this context: one can state falsehoods with clarity and precision. As Carveth Read (1898, p. 351) famously observed, “it is better to be vaguely right than exactly wrong.” A clear and simple benchmark, however precise and wide-ranging, is of little insight if it is arbitrary or otherwise unjustified. The core function baseline therefore cannot

be justified by breadth, precision, and simplicity alone. It must also be evaluated on substantive grounds.

Review of the substance of the core function benchmark suggests that there are powerful reasons to be skeptical that it provides a legitimate baseline against which to measure tax expenditures. Several of these reasons have been raised many times about tax expenditure analysis in general, yet “proponents of tax expenditure analysis generally have failed to respond convincingly to the important criticisms leveled against it” (JCT, 2008, p. 7). The DFC framework seems to provide no exception to this observation. As an initial illustration, recall that Burman characterized the selection between an income and a consumption tax as “the most contentious issue” in defining the normal tax (2003, p. 618). The DFC characterizes the “most fundamental” aspects of the tax system to include a personal and corporate income tax, as well as a Goods and Services Tax. This alone is a controversial normative judgment of enormous significance for the results of the analysis, but no justification is provided. Characterizing the baseline in terms of an income tax then entails defining the tax base. Again, this was one of the primary challenges faced by Surrey, and it has special significance in the Canadian context. Bittker argued that Surrey’s interpretation of the Haig-Simons definition was idiosyncratic and unprincipled. In Canada that definition was the “intellectual inspiration” (Duff, p. 2) of its landmark Royal Commission on Taxation in 1966 (the Carter Commission). The Carter report favoured a comprehensive (Haig-Simons) income tax base, noting that for the purpose of Canada’s tax administration income should be defined as “the sum of the market value of goods and services consumed or given away in a taxation year by the tax unit plus the annual change in the market value of the assets held by the unit” (p. 39).

Haig-Simons is standardly the most fundamental or core definition of income, and it has the virtues of wide-range and simplicity. In addition, both Canadian tax policy and tax expenditure analysis in general have faced criticism for unprincipled variation from that definition. Therefore, one might expect the DFC to adopt the definition or provide convincing reasons for departures from that standard. Instead, the DFC again obscures their inevitable normative judgments in language intended to make the benchmark appear simple, objective, and less subject to interpretation than alternatives. Their report explains that “the benchmark income tax base can be considered a **variant** of the comprehensive income tax base as was first defined by economists Robert M. Haig and Henry C. Simons” (fn. 4, emphasis added). The authors are of course well aware of the implications of adopting a truly comprehensive income tax base:

The comprehensive income tax base would require the taxation of real current additions to purchasing power, or real increases in wealth, which would cover worldwide income from all sources—labour income, rents, dividends, interest and capital gains (adjusted for inflation), transfers, imputed rent on owner-occupied dwellings, the imputed value of household services, and gifts and inheritances. A strict application of the Haig-Simons base would make corporate income tax redundant since income earned at the corporate level would be taxed as it accrues to individuals (fn. 4)

Yet the variant adopted by the benchmark tax system does not characterize the income tax base in this way. It excludes non-market transfers such as gifts, inheritances, and child-support payments. It also excludes the welfare benefits of imputed rents and non-market household services. As it applies to capital, it assumes the benchmark tax system operates on a realization rather than accrual basis, to the enormous benefit of capital holders. It therefore counts the tax-deferral costs of RRSP accounts as an expenditure, while at the same time it excludes from

reporting the cost of income deferral on capital gains.<sup>68</sup> Many of these exclusions are contrary to the notion of income recommended by the Carter Commission, and they are clear departures from the “most fundamental” interpretation of income.

Another problem for the benchmark tax system is its reliance on the status quo. The DFC seems to accept that, in the absence of ‘neutral’ considerations in favour of some characterization of the benchmark, it is appropriate to defer to the status quo. The difficult question of how taxation ought to be structured is avoided with appeal to the way things are. The accepted-in-practice variant of the Haig-Simons definition of income is one implicit concession to the status quo, but there are others. Indeed, sometimes they are made explicit features of the benchmark. Perhaps the best example of the latter type is the benchmark characterization of the basic rate structure. Is the structure of progressive income tax rates part of the benchmark, or is it an expenditure deviation from some other baseline? What is the “most fundamental” or “core” rate structure? The DFC simply accepts the “rate and bracket structure as it exists at any given time.” The credit for the basic personal amount (an estimated ‘cost’ of \$46 billion in 2022) is therefore not an expenditure, simply because it is part of the existing rate structure as a zero rate on some level of income. Similarly, the rate structure for the corporate income tax is also defined as “the statutory general corporate income tax rate in effect at any given time.”

The way (some) things are provides a ready baseline for measuring changes in the system, and indeed it is a distinctive feature of baseline problems that they are often addressed with appeals to the status quo. But such appeals are profoundly misleading without some account of the underlying structure of entitlements. Consider how characterizing the core structure of the

---

<sup>68</sup> Presumably this is only because it would be especially difficult to estimate the liability implied by taxing capital gains on an accrual basis. For example, Henry Simons (1938, p. 207) argued that a realization requirement is a presumption of any realistic income tax and so automatically part of the income tax structure. That strength of practical constraint is not obvious, and whether it is ‘principled’ less obvious still.

tax system in terms of the status quo rate structure obscures the sense in which that very structure distributes enormous economic advantages. Reporting tax expenditures is supposed to reveal how the tax system is used to achieve policy goals for certain targeted groups, such as the clergy in the opening example. But targeted groups is understood rather broadly, as the 3.5 million families claiming the Child Tax Benefit attests. Yet the top federal income tax rate applies to income over about \$215,000, and in 2019 only 287,490 tax filers reported income over \$250,000.<sup>69</sup> The top marginal rate therefore applies to a group of individuals far smaller than other “targeted groups” identified by a tax expenditure analysis. A low federal rate on top incomes is a major benefit for the wealthy, a fact clearly visible if the economic outcome is evaluated against a baseline defensible from a tax policy point of view. Nevertheless, structural advantages reaped by a handful of wealthy Canadians are characterized as fundamental and a core function of taxation, while raising families out of poverty counts as a targeted expenditure. Against a defensible baseline the top marginal rate may be an enormous negative expenditure, while in Canada it is simply built into the benchmark as part of the baseline against which other tax measures are evaluated.

The necessity of an underlying account of entitlements is also starkly visible when reflecting on the non-taxation of personal property (including income) of status Indians and Indian bands situated on reserve. This tax measure is found in Section 87 of the *Indian Act* though the DFC cites insufficient data to estimate its cost. Nevertheless, the undetermined cost is classified as an expenditure—a “non-structural tax measure.” It is not difficult to estimate the ‘cost’ of this measure. Statistics Canada reports that “In 2015, there were approximately 134,000 status Indians living on reserve with market income totalling \$3.3 billion” (Lester, 2019, p. 261).

---

<sup>69</sup> Statistics Canada, High income tax filers in Canada (Longitudinal Administrative Databank)

John Lester estimates that “the distribution of income was such that an average federal tax rate of 6 to 10 percent would be plausible, implying federal taxes ranging from \$185 to \$310 million” (2019, p. 261). Since the DFC lists the non-taxation of that income as an expenditure, that estimate reflects a cost as defined in the DFC report. Yet this implies that it is a fundamental element of the tax system that indigenous people and their income on reserve are subject to federal tax liability, and thus that refrain from such taxation is lost revenue. Note that the DFC is committed to claiming *not only* that Canada is so-entitled (a dubious claim), but the stronger claim that such entitlement is part of a benchmark that is “simpler and less subject to interpretation” than the claims made in so-called “normative” tax policy analysis.

In summary, the DFC relies on a number of controversial normative claims in characterizing the benchmark tax system and then declares its model distinct from normative alternatives. The normative claims I highlighted are the following. It assumes that the most fundamental aspects of the tax system include a personal and corporate income tax, rather than some other tax base. It assumes that the income tax base is defined as a specific variant of a more general concept, and it does not justify the underlying concept nor the variance. It assumes that the fundamental income tax tracks the rate structure as it stands at any given time. It assumes tax jurisdiction over a number of tax units, even when those units are exempt by federal (and even constitutional) law. In addition to these contestable normative assumptions, there is also something peculiar about the counterfactual involved in reporting tax expenditures which leaves it unclear what exactly a tax expenditure estimate amounts to.

### 5.2.7.2 Counterfactuals and Nominal Revenue

Tax expenditures are routinely quantified as ‘foregone revenue’ and ‘costs’ but those descriptions are misleading. Here are some typical comments from the BC Auditor General’s report (2018, emphases added):

[B]ecause government *gives up [tax expenditure] revenue*, it represents a cost in support of policy objectives.

[Tax expenditures are] *foregone revenue*, or money that government doesn’t collect, but could if it didn’t offer that tax break.

In 2016/17, tax expenditures in B.C. accounted for over \$7 billion in *foregone revenue*.

The U.S. General Accounting Office also opens with reference to “[t]he growing *revenues forgone* through tax expenditures reduce the resources available to fund other programs or reduce the deficit” (1994, p. 2, emphasis added). Moreover, the standard practice of reporting tax expenditures as costs faces the same challenge to conceiving of them as ‘revenue foregone.’ This challenge is that both ‘costs’ and ‘revenue foregone’ appear to imply that under a different tax system (namely, one without the expenditure) the value of the expenditure would be, or at least could be, collected or realized. Thus at the heart of a tax expenditure report is a counterfactual claim, a claim about how things would be under some other conditions.

It is difficult to identify what the tax expenditure counterfactual amounts to. Moreover, it is not clear that a meaningful counterfactual can be constructed without specifying the entitlements at stake. Consider an absurd example to make the point. Canada might notice that Alaska has combined state and federal revenues of about \$10 billion. Not taxing those revenues at 25% costs Canada \$2.5 billion in foregone revenue. But plainly not taxing those revenues costs Canada nothing, presumably because Canada is not entitled to tax them. But the DFC and other modern tax expenditure frameworks are reluctant to defend a baseline of entitlements,

preferring to search for ‘neutral’ and ‘objective’ means of specifying the relevant benchmark. They might point out that Canada *could not* tax those revenues. But how should ‘could not’ be understood? Is it a political claim, or a legal claim, or something else? After all, Canada has no political or legal right to the revenues of all Indigenous persons and their income on reserve, yet it claims foregoing taxation of that income as an expenditure.

The counterfactual involved in ‘foregone revenue’ also assumes that everything is held constant except the tax measure under analysis. This is made explicit by the DFC and in other reports. For example, the DFC report indicates that in 2022 the Charitable Donation Tax Credit will provide over \$3 billion in tax reduction benefits for individuals, at the cost of reducing the personal income tax revenue by the same. The implication is that in a counterfactual world without that tax measure the government really would collect \$3 billion in additional tax revenue. Yet, as the DFC report concedes in its discussion of interpretation, the counterfactual is not so straightforward. The first problem is that the modeling that generates the report findings assumes that “the existence of a tax expenditure does not affect taxpayer behaviour” (DFC, 2021, p. 18). But this is almost never true. In the absence of the Charitable Donation Tax Credit, people will look for other ways to minimize the tax implications of their charitable giving, including making use of other tax preferences. Because of the interaction among tax provisions, “the sum of tax expenditures provides very little information about the cost of tax expenditures as a whole to the federal government” (Burman, 2003, p. 615). Economic behaviour is dynamic, and changes in the tax system result in changes in economic behaviour. This is especially clear with tax expenditures designed to attract a certain kind of economic activity. The partial inclusion of net realized capital gains in the income tax base comes at an estimated ‘cost’ of over \$20 billion to the combined personal and corporate income taxes. Yet were this tax expenditure eliminated,

investors would seek other investment vehicles and tax measures to improve their overall economic position.

The dynamic relationship between tax measures and behaviour also includes intergovernmental response. Governments at all levels design their policy in accordance with the background tax structures. Changes in some tax measure often result in corresponding changes in government policy, with reverberations throughout the entire social economic system. For instance, in British Columbia the Persons with Disabilities designation results in monthly payments, supplements, and job support. Of course, the province does not construct this complex social program in isolation from the overall tax-and-benefits landscape. Instead, its policy guidelines and public information directly reference interrelated federal programs including the Disability Tax Credit (a tax expenditure costing about \$1.3 billion) and the related Registered Disability Savings Plan (a tax expenditure costing about \$80 million). These federal initiatives enable a small amount of long-term savings for people with disabilities. As such, the province provides no assistance at all for people with disabilities attempting to save for retirement. In the counterfactual world without those federal tax expenditures, it is plausible that other levels of government would respond to the gap in services, shifting the tax-and-benefits landscape in response to the change in tax expenditures. In short, tax expenditures do not exist in a static environment. They interact with each other, with behaviour, with governmental response, and with overall economic activity (GAO, 1994).

It turns out that a tax expenditure cannot plausibly represent a claim that, in the absence of that tax measure, government should, would, or even could, collect the revenue at issue. So what exactly is being measured? As Burman explains, “[a] tax expenditure estimate reflects the amount by which tax liability is reduced due to a particular tax provision, but it does not measure

the revenue that would be gained by eliminating that provision...” (2003, p. 614). The JCT makes a similar admission, describing expenditures as reductions in “nominal revenues”:

[T]ax expenditure figures calculate the nominal revenues forgone by the existence of the rule in question, not the revenues that would be raised by repealing the rule; the two are not the same because actual repeal would have behavioral consequences that would affect post-repeal revenue collections. Moreover, tax expenditures are not additive, due to behavioral and other issues (2008, p. 5, emphasis added).

So tax expenditure figures are nominal revenues foregone—revenues *in name only*. But then, contrary to its presentation by the BC Auditor General and the DFC, tax expenditure analysis does not reveal ‘costs’ in any meaningful sense.<sup>70</sup> Reports do not measure revenue foregone (revenue the state could or would collect in the absence of the tax measure), and nor do they reveal revenue the government is entitled to collect but does not. Tax expenditure analysis can assess revenue foregone ‘in name only.’ Worse still, that nominal figure is measured from a baseline which is, by design, not normatively defensible. Thus the influential tool of tax expenditure analysis amounts to observing that revenue is distributed one way, which is nominally different from some hypothetical but unrealizable tax liability structure.

Consider how feeble the analysis appears when applied to a real example. The DFC reports that the Disability Tax Credit ‘costs’ the government \$1.3 billion in lost revenue. Does this mean the government is entitled to that \$1.3 billion but does not collect it? It does not, because the ‘core function’ baseline from which the \$1.3 billion is derived is not a specification of entitlements. Does it mean instead that, in the absence of the credit, the government would collect \$1.3 billion in revenue? It does not, since revoking the credit would result in people changing their behaviour, governments changing their policies, and so on. So what does the

---

<sup>70</sup> “Thus, the revenue loss estimate for each tax expenditure is far from precise. It is not an accurate measure of what would be gained by repealing a particular tax expenditure, and, therefore, not an accurate measure of the cost of that tax expenditure” (Fleming & Peroni, 2010, p. 178).

‘cost’ claim amount to? The \$1.3 billion in ‘cost’ is simply a measurement of the difference between a hypothetical liability structure and Canada’s existing tax regime. Yet that hypothetical liability structure does not represent a counterfactual that is normatively defensible or even empirically realizable. The utility of such a report is at minimum unclear.<sup>71</sup>

In summary, the lauded Canadian model relies on a number of fundamental normative judgments in characterizing the normal tax baseline. At the same time, however, its ostensible rejection of a ‘normative’ benchmark leaves it unable to characterize expenditures as deviations from an entitlement baseline. As a result, the DFC analysis is at risk of measuring nothing other than the distance between the Canadian tax system and some arbitrary hypothetical liability structure.

### **5.2.8 The JCT Reform Proposal**

If even Canada’s leading attempt to analyze tax expenditures leaves unresolved the standard criticisms then perhaps tax expenditure analysis needs significant reform. There are basically two ways reform might proceed. The first involves doubling down on neutrality and trying to cleanse the evaluative baseline of as many value judgments as possible. Some normativity is unavoidable, but less is better: “The moment that judgment enters the analysis the potential for internal inconsistency and the appearance of arbitrariness increases. The more judgment that is exercised, the less useful the tax expenditure budget becomes for analytical

---

<sup>71</sup> Fleming and Peroni suggest that even though reported estimates are near meaningless numbers they still give a sense of comparative magnitude and “help dispel the political illusion that tax expenditures are somehow ‘free,’ unlike direct governmental expenditures” as well as “help make the resulting revenue loss a part of budget consciousness” (2010, p. 179). In an earlier work (2008, p. 449) they described reporting as “an important tool for spotlighting hidden government spending and subjecting it to a more rigorous analysis than it would otherwise receive” and they maintain that (p. 499) “an accurate dollar total of tax expenditures is not important. Instead, the critical objective is to ensure that tax provisions with a subsidy or incentive element are carefully examined to see if they are justifiable.”

purposes” (Hall, 1999, p. 8). The second involves accepting that tax expenditure analysis is a thoroughly normative project. Instead of avoiding value judgments, the project would defend them as reasonable or justified, rather than subjective and idiosyncratic.

The JCT and other experts clearly prefer the first approach: “our proposed definition of tax expenditures thus [is] intended to be as transparent and objective as possible” (2008, p. 43). However, in the following I argue that such aspirations for neutrality only serve to camouflage a normative core: “an illusion of value-free scientific precision in a heavily politicized domain” (Kahn & Lehman, 1992, p. 1663). As a result, there is no real choice between neutrality and normativity. Instead, the choice is between a clear expression of guiding value judgments (candid normativity) and various attempts to assert neutrality while denying or obscuring the normative dimensions of the project (cryptonormativity). “Halftruths are often more deceptive than silence” (Bittker, 1969, p. 259) and the JCT reform proposal is a useful illustration of how ambitions of neutrality inevitably reduce a project into cryptonormative obfuscation.

### 5.2.8.1 Perceived Neutrality

In 2008 the JCT set out to transform tax expenditure analysis with its new paradigm. It emphasized perceived neutrality as the key virtue of its proposal, the “overriding objective in rethinking tax expenditures” (2008, p. 42). Critics had long lambasted the normal tax as arbitrary, biased, subjective, and reliant on covert value judgments (e.g. Hall, 1999; Kahn & Lehman, 1992), and the JCT was at least partially sympathetic to that view:

[T]he “normal” tax is largely a commonsense extension (and cleansing) of current tax policies, not a rigorous tax framework developed from first principles. As a result, the “normal” tax cannot be defended from criticism as a series of ultimately idiosyncratic or pragmatic choices. If tax expenditure analysis is to enjoy broad support, it must be seen as *neutral* and *principled*; unfortunately, the “normal” tax

satisfies these requirements only in the eyes of those who already believe that the “normal” tax accurately captures their personal ideal of a tax system (2008, p. 36).

The choice of language in ‘*be seen as*’ is no accident, as they affirm throughout the report. A *neutral* account is repeatedly described as one capable of commanding expert consensus, though not necessarily because it is true, indeed regardless of whether it is true. In a striking comment following their review of the most prominent criticisms of tax expenditure analysis they write,

Whether these criticisms are correct is less important for this purpose than is the observation that they are prevalent. Tax expenditure analysis is useful to policymakers only if its conclusions are widely agreed to be neutral, in both a political and an economic sense (2008, p. 37)

Again and again they insist that the primary reform objective ought to be to “refashion [tax expenditure analysis] in a manner that will generally *be viewed as* more neutral and more principled than the current implementation” (2008, p. 37, emphasis added). As reviewed above, ambitions of perceived neutrality and expert consensus have always been plagued by the myriad conceptions of how the tax baseline ought to be structured. There is nothing close to consensus on how to specify that baseline for comparison, yet “[t]he unavoidable problem is that, by definition, tax expenditure analysis requires comparing actual rules to some hypothetical” (2008, p. 16).<sup>72</sup> This reality has cast the whole project into doubt since the beginning:

[E]very man can create his own set of “tax expenditures,” but it will be no more than his collection of disparities between the income tax law as it is, and as he thinks it ought to be (Bittker, 1969, p. 260)

[A] baseline necessarily reflects ad hoc value judgments, since a value-free ‘correct tax structure’ is impossible (Thuronyi, 1988, p. 1168).

---

<sup>72</sup> One semi-exception is Michael McIntyre’s (1980) proposal that tax expenditure analysis work from a “rhetorical baseline.” His suggestion was that if tax measures are proposed for spending-type benefits then they should be studied and evaluated as expenditures. Thus, the language or rhetoric of tax rules picks out expenditures without appeal to any normal tax baseline. Presumably, advocates would simply change the way they promote their preferred tax measures, i.e., “make the necessary rhetorical shifts” (see Fleming & Peroni, 2010, p. 150).

Even the JCT acknowledged (with evident regret) that “no effort along the lines of a tax expenditure analysis can ever be entirely value-free” (2008, p. 16).

Nevertheless, the guiding idea was to develop a paradigm that made use of a non-normative benchmark for measuring expenditures; it was thought that such a benchmark could provide the necessary baseline but, at the same time, avoid the perils of normative characterizations of the income tax. In what follows I explain how the JCT settled on the conjunction of two baselines that it characterized as neutral and objective, one from law and the other from economics. In short, the legal baseline is characterized in terms of general rules found in present law, and the economic baseline is characterized in terms of efficiency. These two baselines in turn generate two categories of tax expenditures: tax subsidies from the former and tax-induced structural distortions from the latter. I explain the reasoning in favour of these baselines, but I suggest that perceived neutrality is at best a mirage. Both baselines only cloak rather than avoid the crucial and necessary normative judgments grounding the analysis.

#### **5.2.8.2 Cryptonormativity in ‘Present Law’**

The principal advantage of deriving a tax expenditure baseline from existing law is that law does not reflect the idiosyncratic value judgments of any particular individual. Earlier I cited Bittker’s prominent criticism of tax expenditure analysis, which is that there are as many lists of tax expenditures as there are personal conceptions of the ideal tax. But since existing law does not reflect such a personal conception, it ostensibly avoids Bittker’s critique. Thus, the JCT rejected all attempts to characterize the normal tax which are “entirely exogenous to existing law.” Consensus about the ideal tax base will presumably never be achieved, but controversy over the state of existing law is comparatively muted. So their thought was that rather than

becoming mired in value-laden debates about the ideal tax, we might instead look at the existing tax system along with “circumstantial evidence” such as records of legislative intent (2008, p. 16). In this way tax expenditure analysis can bypass debates about, for example, whether the ideal tax base is income or consumption. We need only review the real tax system and observe that the income tax is fundamental to that regime.

Measures in the existing tax regime can then be classified as general rules or exceptions, and the exceptions classified as tax subsidies. A tax subsidy is thus defined as a tax measure “deliberately inconsistent with an identifiable general rule of the present tax law [. . .] that collects less revenue than does the general rule” (2008, p. 39). There is “occasional uncertainty” about the clarity and generality of some tax laws and some judgment is required in determining “what constitutes the general rule, and what the exception.” But such judgments are not particularly concerning, as they are sufficiently distinct from “the kinds of normative judgments that go into the construct of the current ‘normal’ tax.” The advantage is therefore that debates about the content of present law are relatively narrow, technical, and legal rather than broad, vague, and normative.

The ‘present law’ standard for identifying tax subsidies does not achieve the JCT’s ambition of neutrality. In fact, it is cryptonormative because it has not avoided making serious and controversial normative judgments; it has only moved those judgment out of sight. To see this, suppose we grant the debatable assertion that interpretive questions about the general rules constituting present tax law are mostly objectively decidable. We accept, that is, that *if* present law supplies the relevant baseline, *then* we may have mostly technical questions about what the law is. Still, this gives us no reason to affirm that present law *does* supply the baseline. Affirming that antecedent is a crucial normative choice distinct from the normative implications of its

consequent. Reliance on present law is supposed to provide a solution to the normative controversy inherent in defining the normal tax. Yet *if* we accept some ‘normative’ description of the normal tax (e.g., Surrey’s), *then* we may have mostly technical debates about measuring deviations from that tax. Of course, the challenge was to explain *why* tax expenditure analysis should accept one characterization of the normal tax over all the others. Adopting the present law standard is proposed to transcend “sterile debate as to the appropriateness of the choice of [the normal tax] base.” But then simply adopting any arbitrary standard would transcend the debate in the same way. A baseline is not itself ‘neutral’ or ‘objective’ simply because the analysis can proceed ‘objectively’ once the baseline has been adopted.

Imagine a parallel proposal in the domain of moral theory. Controversy persists about how we ought to act. Some say we should maximize well-being, others say we should cultivate good character, still others recommend following a specific set of principles, and so on. Suppose I think a neutral account must find consensus support among holders of these views. But consensus on any of their views is unlikely, so I propose an alternative to circumvent this “sterile debate.” We can derive our standard of action by cataloguing descriptive norms of behaviour, and then analyzing and reporting deviations from those descriptive norms. There may be some disagreements about which behaviours constitute the norms and which the exceptions. But this can mostly be left to the statisticians, and anyway is much less controversial than the broader debate about how people ought to act. With our descriptive norms in hand, we can then make reasonably neutral assessments of when actions deviate from those norms. The central problem for my proposal is that the whole project of measuring deviations from descriptive norms stands in need of justification. The normative controversy was about how to act, and it is unclear how any descriptive norms identified by my project would connect to that question. The ‘present law’

standard for identifying tax expenditures faces a parallel criticism. Tax expenditure analysis is supposed to provide insight into *expenditures*. As we saw at the outset, nothing can be said about expenditures without specifying some structure of entitlements. Adopting the present law baseline simply endorses such a structure (that is, a specific legal one) while masquerading as a neutral choice. But then, like my proposal for discovering what constitutes the moral life, the whole project of measuring from the legal standard stands in need of justification.

### **5.2.8.3 Cryptonormativity in ‘Efficiency’**

The first purportedly neutral baseline for identifying tax expenditures was the general rules of present tax law. The JCT defined ‘tax subsidies’ as deviations from that baseline. The second baseline is efficiency, deviations from which the JCT called tax-induced structural distortions. The motivation to identify a second baseline arose because many long-reported tax expenditures do not satisfy the definition of a tax subsidy, usually because no relevant general rule is clear in the tax code (2008, p. 41). So the definition of tax expenditures as tax subsidies alone was “inappropriately narrow.” At the same time, the JCT was reluctant to expand the scope of their present law standard since that would mean making unprincipled exceptions for specific cases, and so embroil them in the very debates about subjective choices that the whole paradigm was designed to avoid. They viewed the present law standard as objective in comparison with descriptions of the normal tax, and they wanted to hew closely to that commitment. So the category of tax-induced structural distortions does not expand the notion of a tax subsidy, but rather captures expenditures that fall outside its scope. These are “structural elements of the Internal Revenue Code . . . that materially affect economic decisions in a manner that imposes substantial efficiency costs” (2008, p. 41).

Tax-induced structural distortions are therefore tax measures which depart not from legal principles but from economic ones. These distortions impose efficiency costs by reducing the welfare generated through market transactions. However, the JCT acknowledged that such efficiency costs are not all that matters in assessing a tax measure, and an inefficient tax measure might be justified all things considered. But considering all things involves balancing a wide range of considerations, and that is the mandate of the political process (2008, p. 49). They argued that the role of tax expenditure analysis in this process is only to contribute value-free assessments of the efficiency consequences of various tax measures, to “look only to the substantive criterion of efficiency” (p. 42). Whether those consequences are on the whole acceptable is a question outside the scope of tax expenditure analysis.

Identifying and reporting tax-induced structural distortions therefore relies “entirely on an objective inquiry into efficiency considerations” (p. 43). Economists often characterize efficiency as more neutral than other evaluative standards, and the JCT was no exception:<sup>73</sup>

[E]fficiency is an inherently more neutral construct than is equity (and possibly simplicity), and our overriding objective in rethinking tax expenditures is to move to a system that most observers can accept as neutral and principled (p. 42).

The reasoning underlying this assessment is never made explicit in the report, presumably because it is familiar to most readers. The standard observation was articulated in chapter four: a Pareto efficient change makes at least someone better off and no one worse off. The corollary is that an inefficient allocation makes at least someone worse off with no compensatory benefit. Equity, in contrast, involves redistribution from the better off to the worse off. Thus, efficient changes are better for everyone (or at least worse for no one) while equitable changes are better

---

<sup>73</sup> The Joint Committee Chief of Staff later wrote that the “JCT contemplated that tax-induced structural distortions would be analyzed solely under economic efficiency principles, and *not from any normative perspective*” (Kleinbard, 2010, p. 369, emphasis added).

for some and worse for others. As a result, everyone prefers or is indifferent to efficient changes, while equitable changes are more controversial. So, the JCT concluded that concepts like efficiency, equity, and ease of administration “may sound subjective” but “in fact economists have made a great deal of progress over the last several decades in expressing them in objective terms” (p. 48). Efficiency can now be understood in terms of Paretian analysis, equity as horizontal and vertical equity, and so on.<sup>74</sup>

This standard presentation of efficiency characterizes an inefficient state as leaving mutually beneficial gains on the table. The implication is then that efficient changes (Pareto improvements) are just gathering up all those missed opportunities to make someone better off at no cost to anyone. But this commits precisely the error explained in chapter four. It affords overwhelming privilege to efficient improvements over the status quo. The justification provided is that these improvements are uncontroversially preferred by everyone (or at least dispreferred by no one). But, as I argued, that does not make them uncontroversially preferred over Pareto-incomparable alternatives. Crucially, the uncontroversial Pareto principle does not provide *any reason* to prefer Pareto improvements over the status quo when compared with Pareto incomparable alternatives. That preference relies on a separate normative commitment. Again, the “efficient improvement” can always be framed in its explicitly comparative form, revealing the implicit baseline:

- (a) If at least one individual strictly prefers  $x$  to  $y$ , and every individual regards  $x$  to be at least as good as  $y$ , then the society should prefer  $x$  to  $y$ .

---

<sup>74</sup> In case their perspective is insufficiently transparent, they also inform the reader that economists now know that “the ordinary operations of the marketplace do a better job of allocating resources and opportunities than does a more regulated system” (p. 49) and that “competitive markets are likely to improve welfare, by maximizing economic efficiency and productivity” (p. 49).

Satisfying this requirement to specify the comparative baseline makes it nonsensical to describe a tax measure as efficient or inefficient simpliciter. It can only be efficient or inefficient relative to some baseline distribution. In practice, policy analysts tend to drift without explanation between the status quo and some conception of an ideal market as providing the relevant baseline. Consider the ideal standard for tax efficiency articulated by the JCT:

A tax system is perfectly efficient if individuals and firms make the same decisions in the presence of the tax as they would if the tax did not exist, subject only to the fact that they are less wealthy by virtue of paying the tax (2008, p. 52).

This “same decisions” ideal is described by economists as non-distortion and is considered a “major tax policy principle.” It is explicitly endorsed by, for example, the U.S. Treasury Department:

[T]axpayer behavior should be altered as little as possible from what it would be in the absence of a particular tax provision (Fleming & Peroni, 2010).

Notice how the first statement of tax efficiency begins with a “tax system” and then moves to “the tax.” In this second statement the focus is clearly on a “particular provision” rather than the overall system of taxation. But what is the relationship between changes in behaviour resulting from particular tax provisions and lost efficiency? The background assumption is that there exists a counterfactual condition of undistorted preferences which is disrupted by taxation. That condition is an equilibrium in which no changes in the social state can improve welfare without making someone worse off; that is, it is Pareto optimal. Taxes change the market clearing prices in that equilibrium and so reduce the quantity of mutually beneficial transactions. For instance, assuming some price elasticity of supply and demand, a tax on some good raises the cost of supply and so that good is produced less, and the tax also raises the cost to buy, and so it is consumed less. As a result, fewer mutually beneficial transactions occur on the taxed good than would occur without the tax.

Here again we see an application of the argument outlined in chapter four. These missed transactions are described as deadweight loss since they benefit no one, not even the tax collector. Economists also identify deadweight loss in the allocative consequences resulting from changes in market clearing prices. That is, taxpayers adjust their consumption in response to the change in prices. They consume goods lower in their preference rankings and so net less welfare than they would without the tax. For example, they may prefer peanut butter to cream cheese in the no-tax models but because of a tax on peanuts they consume cream cheese instead.

Economists identify these and other ‘distortions’ in all real transactions. Standard examples include that an income tax penalizes delayed consumption (saving) as compared with a consumption tax, that it incentivizes people to work more than they would in a no-tax world to achieve the same income (an income effect), and also that it incentivizes people to work less since leisure is untaxed (a substitution effect). Economists also claim that taxes distort preference ranking, such that a tax on peanut butter incentivizes purchases of cream cheese. Similarly with investments, a tax on corporate earnings combined with a personal tax on the distribution of those earnings is said to distort investment away from corporate equity and towards tax-exempt investment in state and local securities. An example from employment contracts is that the exclusion of health care benefits from taxable income distorts consumer preferences towards additional health care. And so on.

There can be no doubt that taxes change behaviour. But the notion of an *inefficient* ‘distortion’ depends on a more substantial claim about Pareto optimality. Inattention to the relevant baseline can easily mislead, as is evident with the JCT’s example of deadweight loss:

[I]magine that a taxpayer on a limited budget today prefers peanut butter to cream cheese, and a subsidy then is created for production of the latter through the adoption of tax expenditures aimed at cheese makers. In light of the now-lower subsidized price of cream cheese, the taxpayer (and others like her) may choose to switch some

of her peanut butter consumption to cream cheese. Proponents of the cheese tax expenditure would point to the surge in cream cheese production and consumption as demonstrating that an untapped need had been filled, but economists apprised of the facts would conclude that the concomitant decline in peanut butter consumption told a sadder story of deadweight loss, in which peanut butter fanciers distorted their pretax preferences to reduce their tax liabilities.

Note how the conclusion essentially depends on the conception of “pretax preferences.” This state of affairs provides the baseline with which the post-subsidy consumption choices are compared. But what are “pretax preferences”? The pretax preference rankings of peanut butter and cream cheese are but one element in a vast matrix of ‘distortions’ that includes the entire system of taxation and governance, as well as all other market influences classified as ‘distortions.’ The implicitly accepted baseline is therefore the full set of existing distortions less a single one—the cheese maker subsidy. The JCT example thus compares preferences under one enormous set of distortions (pre-subsidy) with another (post-subsidy). But why think changes from the ‘distorted’ pre-subsidy baseline are inefficient? The economists’ conclusion is that there is not merely a change but a *loss in efficiency* when moving from the first cluster of ‘distortions’ to the second. But recall that a loss in efficiency requires at least two alternative social states. If  $x$  is Pareto superior to  $y$ , then everyone either prefers or is neutral to  $x$  over  $y$ . Yet the cheese maker subsidy case supplies no reason to think that everyone prefers the pre-subsidy world to the post-subsidy world. And why would they? Presumably the cream cheese suppliers were worse off pre-subsidy, and the peanut butter suppliers worse off post-subsidy. Comparing to the two worlds reveals that neither is a Pareto improvement over the other: they are simply Pareto incomparable.

The economic reasoning in the sandwich spread example is misleading because it relies on a kind of bait-and-switch. I noted above how the economic notion of deadweight loss is derived from a model environment in which there are no taxes at all. Compared with that “pretax

world,” taxes change the market clearing price and reduce the quantity of goods transacted. Considered in isolation (i.e., ignoring the expenditure side of the budget) this imposes an efficiency loss in that model environment equal to the welfare lost from the missed transactions. Cases like the sandwich spread example also rely on the notion of a “pretax preference” or a quantity of transactions “pretax.” But we must be careful to distinguish between ‘the pretax world’ (the world before this *particular* tax measure) and *the pretax world* (the world with no taxes whatsoever). The former is the baseline in the sandwich spread case, but the latter is the baseline in the deadweight loss demonstration. A reduction in ‘welfare’ in a highly schematic model against a zero-tax baseline does not correspond to the real-world baseline in which all ‘distortions’ are present except the tax measure under consideration (e.g. the cheese-maker subsidy).

To clearly see the bait-and-switch we can suppose the pre-subsidy quantity of cream cheese transactions is 5 (CC5) and peanut butter is 10 (PB10). A tax expenditure then decreases the suppliers’ cost to produce cream cheese, and in turn reduces the price to buyers. This changes the quantity transacted, which is now CC8 and PB2. The JCT explains that this surge in cream cheese popularity is seen by “economists apprised of the facts” as a sad case of deadweight loss—a distortion of pre-subsidy preferences. But now suppose a tax on milk is introduced, and it offsets the effects of the subsidy on cheese production. Quantities transacted return to their pre-subsidy levels. Against the post-subsidy baseline of CC8 and PB2 the return to CC5 and PB10 is a distortion. The absurd result is that CC5 and PB10 is both a ‘distortion’ (post-tax) and not a ‘distortion’ (pre-subsidy), even though both quantities are exactly the same. The only way to avoid this result is to posit that the pre-subsidy preferences are somehow *undistorted*. But that is exactly the switch, since that is only true by assumption in the model environment.

Incentives and subsidies, including tax expenditures, enhance the after-tax economic returns from affected activities and cause taxpayers to engage in those activities to a greater extent than would otherwise be the case, distorting taxpayer behavior and interfering with the free market's allocation of resources (Fleming & Peroni, 2010, p. 164).

The world prior to the cheese-maker subsidy is obviously not an ideal instantiation of “the free market’s allocation of resources.” Indeed, the thought that pre-subsidy preferences are somehow undistorted is inconsistent with the broader conclusion that distortions are ubiquitous. There is simply no reason to accept the pre-subsidy state as the undistorted baseline against which any tax measures must be assessed.

This kind of bait-and-switch is misleading but common. Behavioural changes resulting from taxation are often characterized as inefficient distortions by comparing them with the status quo prior to introducing the specific tax measure under consideration. For example, Fleming and Peroni (2002, p. 143) write that “[t]he denial of deductions for income-producing costs, however, would amount to imposing a tax penalty, thereby causing an economic distortion.” Here are two more examples taken from the JCT report:

Housing tax subsidies can also be viewed as inefficient . . . [because] they encourage private capital to be diverted into the housing sector from other investments that would have been made in a world without such incentives (JCT, 2008, p. 49).

[T]he taxpayer has chosen, in light of the imposition of taxes on his wages, to work less than he would in a world without such taxes (JCT, 2008, p. 53).

These examples plausibly assume that, without those specific tax measures, taxpayers would choose differently. But that is not, in itself, any evidence of inefficiency. A change in taxpayer behaviour is only a distortion relative to some undistorted baseline, but a world without housing tax subsidies or a wage-tax is not such a world. Indeed, no such world is even conceivable. All existing preferences are obviously conditioned by the environment of their holder. The tax-free model environment suggests that in the absence of taxes only undistorted preferences would

remain. Yet the absence of taxes is not well represented as our world with just a few changes, a tidying up of economic distortions. This commits the error identified by Carl Shoup (1969) a half century ago of relying on a counterfactual that “is conceptually invalid because it postulates, for implicit comparison, a state of affairs in which there are no taxes whatever. . .” (p. 577). That world is the state of nature, not a world in which preferences for peanut butter and cream cheese are cleansed of the undue influence of governance:

There can be no other legal subject - no courts, no legislature, no functioning democracy - without a tax system. And as the tax system of a country fractures and loses legitimacy, so does representational government. There is nothing more important, because everything of value that a government does, from national security to healthcare, flows from the tax system (Rosenbloom, 1987, p. 425).

The JCT and others also apparently accept the thesis that efficiency is an approximable continuum, such that fewer tax-induced distortions in market decisions must result in a more efficient outcome. The idea is that a perfectly efficient tax is distinguished by its total absence of effect on behaviour, so a tax with less effect must be more efficient.

A tax system is perfectly efficient if individuals and firms make the same decisions in the presence of the tax as they would if the tax did not exist, subject only to the fact that they are less wealthy by virtue of paying the tax (2008, p. 52).

Again, the idea is then that if *no-tax* markets *maximize* efficiency, then *fewer taxes* will result in *increased* efficiency. But this logic has been rightly criticized as mistaken about ‘distortions’ outside of a perfect model environment. Most relevant, the failure of approximation was formally demonstrated to apply to general equilibrium models by Lipsey & Lancaster in their celebrated 1956 essay *The General Theory of the Second Best*. Given that full satisfaction of the model conditions cannot be achieved, the model optimum does not motivate greater satisfaction of the model conditions. In other words, the fact that a tax-free model is efficient when it has no

tax induced distortions provides no reason to think that reducing taxes in the real world will increase efficiency.

In summary, the 2008 JCT recommended increasing perceived neutrality by relying on two objective characterizations of the normal tax baseline: general rules in present law and economic efficiency.<sup>75</sup> They claimed that economists can provide an objective efficiency assessment, and then society can determine how to balance that loss in welfare against other values. Interpreting the present law baseline is probably more controversial than they allowed, but more crucially, the possibility of a neutral interpretation of the baseline does not, in itself, provide a reason to accept that baseline. The purportedly complementary standard is efficiency, but the problem with that evaluative standard is that it either relies on the status quo as a baseline (which requires justification) or it relies on a non-approximable economic model in which case the results do not translate to real tax systems in the straightforward way suggested by the JCT. Those with an “overriding” concern with *perceived* neutrality (even over correctness) may not be concerned. Certainly ‘present law’ and ‘efficiency’ command a great deal of expert allegiance, and both have the veneer of objectivity. But for the rest, the emphasis on neutrality seems unlikely to succeed in delivering a compelling vision for the reform of tax expenditure analysis.

### 5.2.9 Candid Normativity

Tax expenditure reporting is a paradigm case of tax policy work plagued by a baseline problem. Attempts to cleanse the project of a normative baseline have failed. Yet insistence on

---

<sup>75</sup> One interesting feature of the JCT’s renunciation of the normal tax in favour of tax expenditure analysis “that, to some extent, had no normative baseline” is how quickly it was abandoned (Fleming & Peroni, 2010, p. 147). After only two years “the Staff totally reversed itself by abandoning its non-normative approach and reembracing the [Haig-Simons] baseline” (p. 147). Yet debate persists and the Chief of Staff who oversaw the turn to neutrality has even recommended Congress enact legislation adopting the 2008 baseline.

the ideals of value-free neutrality and objectivity remain the norm. I showed above how various strategies are used in pursuit of this norm. However, I also argued that these strategies only serve to obfuscate the values at work in selecting models of the tax structure. For example, citing “efficiency” and “present law” simply defers to the values inherent in those standards. Moreover, the selection of standards itself is obviously an indispensable normative choice. Such choices are then contrasted with “normative” alternatives, as though there are non-normative options available. These and related strategies are misleading—half truths at best.

On this point I earlier cited Bittker’s observation that, in the world of tax expenditures, “[h]alftruths are often more deceptive than silence” (1969, p. 259). Better to say nothing about the values animating the project than to adopt some deceptive fiction about objectively determinable “core functions” or “structural provisions” of the tax system. Probably Bittker is right about that, but what his observation leaves out is also illuminating. In addition to silence and deception there is, I suggest, a third way. What I have in mind is the identification, articulation, and justification of the political value structure on which tax expenditure analysis unavoidably depends. The tax expenditure project is so deeply mired in ambitions of “perceived neutrality” and value-free analysis that, somewhere between silence and deception, expressing the truth has been completely displaced as a viable alternative.

What I mean by the “truth” in this context is not a definitive account of the correct baseline tax structure against which expenditures can be measured. Instead, the truth I have in mind is an *accounting* of the value-laden assumptions inherent in any given model of the baseline tax structure. This requires a serious and sustained project of inquiry into, for example, Canada’s reliance on a Benchmark Tax System. Again, the result would not be a finding of the “correct” baseline tax structure. It would be a clear articulation of the values that ground various

baseline options. If the effort poured into *obscuring* these values was redirected into identifying and defending them, the result would exist in the third way between silence and deception. This is what I mean by candid normativity.

Candid normativity might involve generating several models against which tax expenditures are measured. For example, a tax structure generated by an optimal tax model might have one accounting of expenditures, and that accounting should be accompanied by an explanation of the value assumptions that drive the model. But that should be compared with other models of the baseline structure. It has never been taken seriously in this context, but there may be space for public input into an ideal model. This could democratically ground a model of the tax structure—people could be asked for their ideas of a fair tax system, and expenditures could be calculated from some composite model of responses. Why not generate another model based on expert input from outside economics? Economics is value driven, but the animating values are often left unexplored. As Slemrod and Bakija acknowledge, fairness is essential yet it “is not in the end a question of economics” (2017, p. 60). Or recall the characterization of economics in policy making offered in the introduction of this chapter:

One might without exaggeration say that economics is today the science of administration and policy-making – that every weighing of costs and benefits **in pursuit of the greatest good** (which is what policy-makers are expected to do) is either a form of economic analysis or an ad hoc judgment with no discipline at all (Kronman, 2017, p. 222-223, emphasis added).

Economists work with these ideals, but immense complexity is buried just in the notion of “the greatest good” (including how costs and benefits are conceived and compared). Such ideals need to be dragged into the light and investigated rather than assumed at the outset and then obscured as much as possible in the analysis.

Candid normativity accepts, as do tax expenditure enthusiasts when pressed, that a value-free baseline is impossible. It differs in how it responds to that reality. Tax expenditure theorists tend to push for as little acknowledgment of the animating values as possible. But why? The JCT argued that “if tax expenditure analysis is to enjoy broad support, it must be seen as *neutral* and *principled*.” It is probably true that tax expenditure analysis would lose its utility if it devolved into a multitude of idiosyncratic comparisons of the tax system to individual authors’ notions of an ideal tax system. But this is probably an overblown concern. Any political question—whether about political organization (e.g. voting), distributive justice, constitutional rights, and so on—has abundant value-laden answers. Nevertheless, popular conceptions tend to roughly coalesce around a smaller number of carefully studied alternatives. The same could be true about some limited expressions of the baseline tax system, and the result would be a smaller range of models. Instead of identifying one model as “neutral” or ending up with innumerable different options, analysts could design several. The added effect would be a clear demonstration of how crucial the value assumptions of such models are to their conclusions. It is true that any one model could no longer parade (falsely) as neutral, but the overall process could be seen as *more* neutral due to abandoning that very fiction. In addition, the models would be expressly principled, and the relevant principles articulated and defended.

In summary, tax expenditure analysis insists on perceived neutrality. But that perception is always deluded, because, when pressed, everyone accept that tax expenditure analysis is not neutral at all. This has led analysts to a choice: remain silent about the underlying value assumptions, or obfuscate them with half truths about core functions, efficiency, present law, or other normative standards. My suggestion is that the project instead focus on articulating and defending the value judgments that underlie various tax expenditure models. That is the essence

of candid normativity. It is true that universal consent will never be achieved, but so it goes in all domains of political judgment and discourse. Why should tax be any different?

The tax expenditure case study is intended to illustrate how a theoretical baseline problem can play out in a public policy context. Yet the kind of baseline problem I identify in this project also has much wider reach. The public discourse about taxation is a case in point. It shows how the pretax entitlement baseline is the foundation for by far the most impactful public ideal about taxation: the “fair share.” It is an ideal so widespread that we can hardly think of taxation without it coming to mind. Nevertheless, I argue that it may be worth considering how deeply it too reflects an indefensible baseline.

### 5.3 Taxation and ‘Fair Shares’

A single idea dominates Canada’s public discourse about justice in taxation: the fair share. Perhaps that alone is reason to find it suspect, if enmity is the “thunder around everything that shines.”<sup>76</sup> It is banal to point out that that politics makes strange bedfellows, but the fair share has apparently made bedfellows of us all. Consider a recent parliamentary motion brought by New Democratic Party leader Jagmeet Singh, as the federal government grapples with soaring costs of living in the wake of a global pandemic and Russia’s invasion of Ukraine. Singh urged the government to extend a 3% surtax on “excess profits” of the banking and insurance industries to also include oil companies and big-box stores. His reasoning was extensive, but wealth inequality was “the heart of the matter.”<sup>77</sup> His stated aim was to ensure that everyone contributes

---

<sup>76</sup> Victor Hugo (1887), *Choses veus*, “You have enemies? Why, it is the story of every man who has done a great deed or created a new idea. It is the cloud which thunders around everything that shines. Fame must have enemies, as light must have gnats. Do not bother yourself about it; disdain. Keep your mind serene as you keep your life clear.” A massively influential idea with so little opposition is, perhaps for that very reason, a suspect idea.

<sup>77</sup> For the full discussion see records of the 44<sup>th</sup> Parliament, 1<sup>st</sup> session, 42<sup>nd</sup> seating, March 21, 2022: <https://www.ourcommons.ca/DocumentViewer/en/44-1/house/sitting-42/hansard>

their “fair share.” The importance of this objective was echoed by numerous supporters, each dutifully repeating the fair share mantra.<sup>78</sup>

Yet the basic idea of paying a fair share is impressively non-partisan. Conservative rhetoric also energetically displays a commitment to the fair share principle. For instance, when the *Online Streaming Act* was recently up for discussion, the stated Conservative agenda was to ensure “that the major foreign streamers invest in Canada and pay their fair share.”<sup>79</sup> Party lines were not drawn around the fair share principle, with some parties for and others against. Instead, the disagreement was merely over how to achieve the fair share ideal: regulatory control on the one side, and “consultation and co-operation with the industry” on the other. Moreover, Conservative comments about the *Online Streaming Act* are no anomaly; they are representative of a more widespread commitment to the fair share principle.<sup>80</sup> Parliamentary debate records reveal how their “side of the House” has been “really clear” on precisely this commitment: all economic actors, foreign or domestic, “must pay their fair share in Canada.”<sup>81</sup> As one

---

<sup>78</sup> Jagmeet Singh (NDP) “the ultrarich do not pay their **fair share**”

Daniel Blaikie (NDP) “there are people who are avoiding paying their **fair share**”

Lindsay Mathysen (NDP) “They already do not pay their **fair share** in taxes”

Terry Beech (Lib) “[We have] a firm commitment to ensuring that everyone pays their **fair share** of taxes”

Kevin Lamoureux (Lib) “We want to ensure that everyone pays their **fair share**”

Laurel Collins (NDP) “Are we going to make big banks, big box stores and big oil companies pay their **fair share**?”

Lisa Marie Barron (NDP) “[Access to basic supports] starts with ensuring that everyone is paying their **fair share**”

Rachel Blaney (NDP) “they need to step up and pay their **fair share**”

Francesco Sorbara (Lib) It was “the right thing to do” to “make sure that everyone paid their **fair share**”

Julie Dzerowicz (Lib) “We will continue to . . . ensure that large, profitable banks and insurers pay their **fair share**”

Leah Gazan (NDP) “the ultrawealthy and multinational corporations [must] pay their **fair share**”

Niki Ashton (NDP) “companies [and their beneficiaries] are . . . sheltered by Canada from paying their **fair share**”

<sup>79</sup> <https://www.ourcommons.ca/DocumentViewer/en/44-1/house/sitting-87/hansard>

<sup>80</sup> Cathay Wagantall (CPC) “I totally agree with having the big providers pay their **fair share**”

Glen Motz (CPC) “Large foreign streamers should pay their **fair share**”

G rard Deltell (CPC) “We have to pay our **fair share**”

Matt Jeneroux (CPC) “Foreign streamers should pay their **fair share**”

Don Albas (CPC) CRA is making deals with “big businesses so they do not pay their **fair share** of tax”

Hon. Ed Fast (CPC) “We do believe that the wealthy should pay their **fair share**”

<sup>81</sup> <https://www.ourcommons.ca/DocumentViewer/en/44-1/house/sitting-69/hansard>

Conservative member put it “I understand why some of the members across the way say that everyone should pay their fair share, and we agree with them.”<sup>82</sup>

The fair share is by no means confined to the halls of parliament. During the summer of 2021, Abacus Data was commissioned to conduct a rigorous national public opinion survey about tax fairness in Canada.<sup>83</sup> Their findings were all intriguing; some were astonishing. Perhaps least surprising, they found that only 14% of Canadians think the tax system is fair, and over 70% believe that large corporations and wealthy individuals “pay less than their fair share.” Of special interest is that most Canadians (64%) think the Trudeau Liberals have done “not enough” to ensure everyone “pays their fair share in taxes.” What might be done? An astounding 89% of Canadians support an annual wealth tax and an increase of the top marginal income tax rates. Moreover, the fair share ideal promises real political allegiance. A full 89% of Canadians would either “definitely” or “probably” consider voting for a political leader who “promised to take real and concrete action in making sure everyone pays their fair share and increasing taxes paid by the wealthiest Canadians and large, profitable corporations.” But where should Canadians turn, when every party leader and member alike endorses the fair share ideal? Here, as one might predict, results are far less decisive. When it comes to assessing which party leader would best implement the fair share ideal, 21% favoured Trudeau (Lib), 21% Singh (NDP), 15% O’Toole (CPC), 6% another party, and 36% were simply unsure.

In addition to political rhetoric and public opinion, both governmental and non-governmental tax reform campaigns tend to make ‘paying a fair share’ their central platform. On the side of government, the 2022 federal budget dedicates its ninth chapter to tax fairness. It begins as expected: “Since 2015, the federal government has worked to ensure that the wealthiest

---

<sup>82</sup> <https://www.ourcommons.ca/DocumentViewer/en/44-1/house/sitting-69/hansard>

<sup>83</sup> For the full methodology and results see [https://www.broadbentinstitute.ca/tax\\_fairness\\_survey](https://www.broadbentinstitute.ca/tax_fairness_survey)

people and businesses pay their fair share.” All assessments of fairness are relative to a legal description of pretax holdings. For example,

“The federal government has taken significant steps to increase the fairness of the tax system, including by increasing taxes on the wealthiest one per cent of Canadians. However, some high-income Canadians still pay relatively little in personal income tax (PIT) as a share of their income—28 per cent of filers with gross income above \$400,000 pay an average federal PIT rate of 15 per cent or less, which is less than some middle class Canadians pay.”

On the NGO side, the Professional Institute of the Public Service of Canada represents (among others) about 12,000 tax professionals at Canada Revenue Agency. Their campaign is inspired by a single moral proposition: “Wealthy corporations and the ultra-rich don’t pay their fair share.”<sup>84</sup> Another major national campaign to ‘Tax the Rich’ counts Oxfam Canada, Greenpeace, and the Canadian Union of Public Employees among its partners.<sup>85</sup> Their bedrock commitment is by now routine: “It’s time for the super wealthy to pay their fair share of taxes.”<sup>86</sup>

Why is there so much consensus about the fair share ideal? According to Victor Hugo’s maxim, a truly bright idea ought to attract not a little thunder. Yet the fair share soars through clear skies. In part, remarkable consensus about the ideal combined with ambivalence about its direct implications is simply a reflection of its generality. Perhaps the fair share is what Rawls (1971, p. 5) called a concept of which there are numerous conceptions. After all, what is a fair share? As we’ve seen, one influential idea is that people should pay for what they consume. For centuries people have thought about taxation in terms of “give-and-take” (Seligman, 1908, p. 159). If I use one litre of gasoline and you charge me for two, that’s unfair. I should only be charged for what I use—the “user pays” principle. This is the basis of James Buchanan’s (1976)

---

<sup>84</sup> <https://pipsc.ca/news-issues/tax-fairness>

<sup>85</sup> The full list includes Acorn Canada, Broadbent Institute, CUPE, Greenpeace, LeadNow, Oxfam Québec, Oxfam Canada, Resource Movement, Canadians for Tax Fairness, and United Steelworkers of Canada.

<sup>86</sup> [https://www.broadbentinstitute.ca/tax\\_the\\_rich](https://www.broadbentinstitute.ca/tax_the_rich)

famous fiscal exchange theory of taxation, and the main intuition underlying the principle that taxes paid should be proportionate to benefits received. But many people see important differences between society and a market exchange. Society is a collective undertaking, without which we would have almost nothing at all. Millennia of cooperation now supply the political, legal, cultural, technical, intellectual, and other preconditions that make enormous wealth possible. Perhaps everyone should contribute equally to this common project. But as has been reiterated from the Synoptic Gospels<sup>87</sup> to John Stuart Mill (1848), an equal contribution depends on how much one has. This line of thought grounds the general principle that taxes should be levied in proportion to ability to pay.

So the fair share principle is open to (one might say vulnerable to) interpretation. In the world of politics, it is often pointed out by Conservatives that in nominal quantities corporations and wealthy individuals ‘pay’ a lot in taxes. Moreover, those payments make up a large portion of tax revenue. So the rich earn a lot, but they pay a lot, isn’t that fair? For example, Glen Motz (CPC) remarked that “We always hear from the NDP-Liberal coalition that those corporations do not pay their fair share of taxes, but people should look at what they actually pay to the government in taxes.”<sup>88</sup> Motz does not dispute the fair share principle, but only implies that it is satisfied. Similarly, though with considerably more bluster, John Brassard (CPC) said “it is an absolutely absurd assertion [that taxes on banks should be increased]. Banks pay their fair share. Businesses pay their fair share. Individuals pay their fair share. I just did my taxes, and believe me, I am paying my fair share.” Again, the principle is accepted, but its interpretation disputed.

---

<sup>87</sup> (Mark 12:41-44; Luke 21:1-4) He sat down opposite the treasury and observed how the crowd put money into the treasury. Many rich people put in large sums. A poor widow also came and put in two small coins worth a few cents. Calling his disciples to himself, he said to them, ‘Amen, I say to you, this poor widow put in more than all the other contributors to the treasury. For they have all contributed from their surplus wealth, but she, from her poverty, has contributed all she had, her whole livelihood.’

<sup>88</sup> <https://www.ourcommons.ca/DocumentViewer/en/44-1/house/sitting-82/hansard>

One way to approach the fair share discourse is therefore to think it has the right core idea but needs some theoretical maintenance. We've seen various theories that try and capture and systematize the relevant intuitions about fairness. On this approach, we should set about figuring out whether fair shares should be measured in nominal terms, as implied by Motz, or rather in terms of sacrifice, as suggested by both Mill and the Gospels. Figuring out how much it is fair for people to pay, according to this strategy, might involve a complex balancing of the resources they have access to, how much they have benefited from services provided by the state, as well as general considerations of administration and enforcement. The research program involves taking the fair share as a starting point and then working out better arguments, increased conceptual clarity, and extensive technical modeling. This is more or less the approach taken by modern tax policy analysis (Duff, 2018).

But there is also an alternative. Instead of analyzing and reinforcing the intuitions that guide broad convergence on the fair share ideal, that ideal might be abandoned or at least seriously modified. It is a surprising suggestion, maybe an outrageous one, but I think it is worth considering. First, there is Victor Hugo's observation about good ideas and their enemies. Perhaps the 'fair share' draws no thunder because it is so general that it has essentially no direct policy implications. It becomes an empty husk into which even diametrically opposed policy prescriptions might be tightly packed. It invites multitudinous political commitments precisely because in itself it is devoid of any meaningful content. Still, that alone is not enough to abandon it, at least if we want to avoid consigning the famously ambiguous 'justice' itself to the flames.

The real threat posed by the fair share discourse is that it may confuse the moral stakes by focusing tax justice on a distribution of *payment* rather than *entitlement*. In the first place it may concede far too much to the status quo. It accepts that people have some entitlement to their

holdings and then argues that wealth inequality or some other consideration justifies a reduction in those entitlements. But notice how much more difficult it is to explain why someone ought to have  $x$  than to explain why, given that they have  $x$ , it should not be taken away. Framing the question in the latter way makes things too easy for the Grotian transformation of fact to right. On this point consider the user-pay principle referred to above. The principle assumes that people and companies have bundles of resources with which to pay for services. Yet in a modern, complex economy it is hard to imagine how entitlement to that bundle of resources can be defined prior to the application of fair rules of taxation. So, as has been emphasized repeatedly throughout this project, the user pay principle relies on a given baseline of entitlements in assessing fair taxation, when fair taxation must be determined before any such baseline is conceivable. Nevertheless, the user pay principle invites the Grotian transformation, taking a factual distribution of legal holdings and converting them to *right*.

This suggests that the problem is not with fair shares *per se*, but rather with conceiving of the commitment in the standard way: as *paying* a fair share. We might instead say that people are entitled to *receive* a fair share, but putting the point that way means looking at the problem of fair taxation the other way around. It means looking at what people have and asking whether that reflects what they are entitled to. On this view, the idea of moral entitlements to some legal distribution prior to taxation, and consequently of ‘payment’ out of those entitlements, is simply set aside. But this appears to raise a serious problem. Don’t people care very much about their taxes paid? Consider Tony Judt’s attempt to encapsulate the American mind: “we do not want the state interfering in our affairs. And above all, we do not wish to pay more taxes.”<sup>89</sup> If paying taxes matters ‘above all,’ then how can such an important consideration be simply ignored?

---

<sup>89</sup> <https://www.nybooks.com/articles/2009/12/17/what-is-living-and-what-is-dead-in-social-democrac/>

Quite easily if, as I suspect, the common observation that people care about how much they pay in taxes is highly misleading. Strictly speaking, few (if any) people really care about reducing their taxes paid. To see why this might be plausible, consider *competing for coders*.

*Competing for coders:*

A company named Big Tech aims to lure coders working in another tax jurisdiction. Those coders are currently salaried at \$100,000 with an effective tax rate of about 22%, so they take home about \$78,000. Executives at Big Tech come up with an ingenious plan: they will offer the coders a salary of \$90,000 *but* will emphasize the comparatively favourable tax conditions of their jurisdiction. Salaries in Big Tech's jurisdiction are subject to a lower effective tax rate of about 13.3%, so the coders would take home about \$78,000. However, the executives stress that the decrease in effective tax rate from 22% to 13.3% means the coders will pay only \$12,000 in taxes instead of \$22,000, a reduction of their tax burden by almost half! This is a very significant reduction in taxes paid. If the coders care deeply about reducing their taxes paid then, all else equal, it seems they ought to prefer the new position. Indeed, if they care *at all* about reducing their taxes paid, then the new position offers them something valuable to that same extent. Instead, I think they would find themselves perplexed by Big Tech's offer—why should they care about the reduced tax rate when the after-tax result is the same?

Competing for coders illustrates how so much attention to gross income rests on a dubious assumption. People care about gross income because they see it as connected in a particular way to after-tax income. They assume, that is, that a decrease in the tax rate applied to gross income will directly result in more after-tax income. For instance, they believe the market allocated them a salary which is set out in their employment contract and, in the absence of taxation, they would get all of that amount. Or, at least, with less taxation they would get more of

it. Yet this is just what Carl Shoup (1969) explained “is conceptually invalid because it postulates, for implicit comparison, a state of affairs in which there are no taxes whatever. . .” (p. 577). Besides, on reflection their complaint about taxation is not about their ‘tax burden’ at all. It is rather about a perceived reduction in after-tax resources, which they see as necessarily connected to taxation.

Gross and net income are not static figures in some kind of necessarily inverse relationship. This can be seen clearly by extending the thought experiment to its logical limit. Assume tax burdens are reduced to zero. Under this assumption, there is no distinction between gross-of-tax and net-of-tax income, so individuals receive their full market allocation. The common thought is that as taxes decrease, net income necessarily increases. Therefore, as taxes approach zero, net income approaches some pure tax-free market figure. Unfortunately, the magnitudes are not connected in that way. The no-tax allocation is indeterminate and probably close to zero since there is no society at all. This is what Carl Shoup had in mind when he explained that the hypothetical “is conceptually invalid because it postulates, for implicit comparison, a state of affairs in which there are no taxes whatever, and no government borrowing or creation of new money, hence impliedly no government services, not even of the minimum type and amount necessary to assure existence of the society. . . . [This] objection is conclusive” (1969, p. 577–578). So my suggestion is that people care about what they ‘pay’ in taxes only instrumentally—and only because they wrongly assume that what they ‘pay’ is tightly connected with what really matters to them: the resources they are entitled to after tax.

The ‘fair shares’ discourse therefore suffers from two problems. First, the very idea of taxes paid is a kind of Grotian metamorphosis of fact into right. Second, people care about paying a fair share of taxes only because they wrongly assume that a dollar less in taxes means a

dollar more in their pockets. This suggests that what really matters is the level of resources to which individuals are entitled. If this is right, then perhaps there is some motivation to abandon the fair share discourse after all. Or, at minimum, the discourse might address fair shares of *holdings* rather than insisting that all Canadians *pay* a fair share.

Conceiving of property in terms of after-tax entitlements has the added advantage of encouraging a different perspective on the relationship between ‘taxpayers’ and the state. I think Benjamin Franklin (1789) had something like this perspective in mind when he wrote that “contributions to the public exigencies are not to be considered as conferring a benefit on the public, entitling the contributors to the distinctions of honor and power, but as the return of an obligation previously received.” Wealthy individuals and companies are not doing the state and its citizens a favour by ‘paying’ their taxes, the law is returning possessions to which those entities are not entitled. Assessing whether they have more or less than they ought to have does not involve measuring the distance between their current holdings and some pretax level of holdings. Therefore, the magnitude of that distance provides no basis for celebrating (or condemning) some extent of ‘contributions’ to the public.

With that in mind, let us return to where we started. Again, Singh (NDP) presses the government to divert resources from private industry to the general public.

Mr. Speaker, Canadians are paying \$2 or more per litre at the pumps. People are hurting. Families are hurting. At the same time, big oil and gas companies are posting huge profits, in some cases record profits. We also know that the government has seen an increase in revenues. Will the Prime Minister take the side of people, tax the excess profits of these companies and reinvest that into the pockets of Canadians by increasing the child benefit and the GST credit by \$500 each?

Trudeau (Prime Minister, Lib) replies, in relevant part:

On the issue of asking those who have succeeded most to pay more, we have done that . . . Furthermore, in our latest budget, we are moving forward and asking the

largest companies and banks to pay a little more to help Canadians. That is exactly what we are doing.

For Trudeau, the largest companies and banks have the resources, and the government can only ‘ask’ whether those companies might ‘pay a little more’ of their wealth ‘to help Canadians.’ This reflects a general perspective accepted by the Liberal party. The NDP’s platform for the 2021 federal election contained a fairly modest set of economically progressive policy prescriptions. For example, they wanted to raise the top marginal income tax rate by 2%, introduce a 1% wealth tax, reduce the capital gains exemption, and work to close off-shore tax havens. How did the Liberals respond? The problem, said Trudeau, is that you can’t make the rich pay for everything:

The idea that you can go with unlimited zeal against the successful and wealthy in this country to pay for everything else is an idea that reaches its limits at one point. And I don’t think the NDP get that.<sup>90</sup>

All of this makes sense within the standard fair share discourse outlined above. That discourse suggests we should examine how much wealthy individuals and companies have ‘paid’ in taxes and see if their payments meet some standard of fairness.<sup>91</sup> If so, the most the government can do is ask if the rich might nevertheless contribute a little to help the Canadian public.

The alternative prospect, if defensible, would simply ask what overall bundle of resources big banks, oil and gas companies, insurance agencies, wealthy individuals, and even ordinary people are entitled to. The answer might depend on choices they made, including risks assumed and other measures of personal responsibility, and it might also depend on aggregate welfare, or any number of other considerations. What it would not depend on, however, is how

---

<sup>90</sup> <https://www.nationalobserver.com/2021/09/17/news/election-2021-who-will-fight-tax-fairness>

<sup>91</sup> Trudeau might also be alluding to the risk of capital flight here, but opportunities for avoidance are themselves a policy choice (e.g., Piketty, Saez, & Stantcheva, 2011)

much they ‘paid’ in taxes. This is the basic implication of moving from a principle that everyone ‘pay their fair share’ to a principle that everyone ‘*have* their fair share.’

#### **5.4 Conclusion**

The purpose of this brief accounting of Canadian public discourse about taxation is twofold. In the first instance it reveals just how entrenched the notion of “paying a fair share” is in the public mind. In essence, no matter one’s politics, there is simply no alternative framing of tax justice on offer. I suggest that there is another way to think about taxation: as the mechanism that ensures everyone *receives* a fair share rather than *pays* one. On my view, we begin by asking what grounds entitlements to resources, and what that justification implies for the existing distribution of resources. Distributive justice, in other words, is what determines entitlements. Taxation is indeed one tool for realizing distributive justice, and in doing so it is not constrained by what people possess, but only by what they are entitled to. As this project has emphasized throughout, there is no reason to accept gross income or other pretax holdings as offering a morally salient distributive baseline.

#### **5.5 Chapter 5 Conclusion**

The purpose of this chapter was to identify and explore the theoretically characterized baseline problem in some real world, practical cases. The two cases I chose were the theory and practice of tax expenditures, and the public discourse about taxation.

Tax expenditure analysis emerged as a means of revealing how government uses the tax system to allocate economic benefits outside the ordinary budgeting process: government can “spend” by not taxing. Critics observed that government spending and not taxing are only

equivalent when measured from a baseline of total government entitlement. However, tax expenditure analysis does not accept that government is entitled to “the last penny.” Instead, it maintains that government is entitled to the extent of the normative tax: some baseline description of the tax system. However, the notion of a normative tax is objectionable to analysts keen to persuade outsiders that their work is neutral, objective, and value free. Their response was to emphasize neutrality as much as possible by deferring to value judgments in other domains, obscuring their own value judgments, and contrasting their approaches with “more normative” alternatives. I argued that these attempts are not successful in purging tax expenditure analysis of its normative baseline.

Tax expenditure analysis should abandon ambitions of value-free neutrality and instead focus on articulating and defending its unavoidable link to political judgment. I call this project candid normativity and suggest that it can remain both neutral and principled. It is neutral in the sense that it remains open to various models that are grounded in defensible political values. It does not assert that a single model is a “neutral” description of core or structural functions of the tax system. It is principled in the sense that models are explicitly driven by principled value assumptions. Rather than obscuring these principles, they are acknowledged at the outset. This, I suggest, is an unexplored third way between silence and half truths.

The second case I discussed was Canadian public discourse about taxation, which is dominated by the ideal of everyone paying a fair share. I noted how this ideal is accepted across the whole political spectrum and also suggested this may be reason to approach it with some skepticism. However, I accept that the fair share (like “justice”) may provide a useful concept of which there are simply many competing conceptions. If so, then I suggest a novel conception: *receiving* a fair share rather than *paying* one. This reinterpretation moves the focus from taxes

taking away from some entitlement baseline to taxes functioning as a distributive mechanism in determining economic outcomes. It follows naturally as an attempt to avoid the baseline problem discussed most thoroughly in chapter three. It recognizes that pretax holdings do not reflect a deep moral structure. As argued, because the moral salience of such legal descriptions depends on antecedent justice in taxation, they are not a constraint on securing that very justice. In a modern society, economic justice comes after tax, not before.

## Project Conclusion

This project began with the orthodox conception of tax justice as requiring a fair and efficient distribution of the costs of civil society across the resources owned by its population. From theory to practice, it is clear that this orthodoxy continues to exert significant influence over tax justice. It is less clear, however, that its proponents can adequately respond to the few detractors who, over the past half-century, have argued that the entire paradigm rests on an error. That error is the acceptance of a pretax legal description of holdings as a morally relevant baseline against which to measure and evaluate the impacts of taxation. Rejecting that baseline is a heterodox act, and a direct challenge to the “quasi-religious theology of markets, inequality, and private property” (Piketty, 2020, p. 36).

For that reason, it is not surprising that fealty to the moral significance of pretax descriptions remains the norm. Defenders of such baselines argue that their critics are bound to an implausible conception of property as purely conventional and legal. But property, they point out, is much more than mere convention. However, my argument is that one can accept the latter point while rejecting the former. Property is more than mere convention, yet the orthodox paradigm is nevertheless plagued with a baseline problem. This is because the orthodox paradigm does not provide some independent moral account of entitlement to resources equivalent to pretax holdings. Instead, it takes the legal description of pretax holdings as a moral starting point. That legal description is, if nothing else, a legal convention. As a legal convention, it is but one step in the complex process of calculating legal entitlements. It reflects all the ambiguities, idiosyncrasies, and contingencies of the law, especially the law of property. Any moral salience *in virtue of* its status as a legal determination must emerge from the moral status of the system, including taxation, which is its genesis. Legal descriptions of pretax holdings

therefore cannot stand as fixed moral points that constrain taxation, as their moral salience depends on justice in taxation already being achieved. I maintain that a number of concepts routinely relied on in the practice of tax justice unavoidably depend on the moral significance of legal descriptions of pretax holdings. This includes horizontal equity and vertical equity, as well as their animating ideals of ability to pay (both endowment and equal sacrifice variants) and the benefit principle.

Not only can the baseline criticism of the orthodox paradigm be salvaged, but I also argue it can be extended. It has been applied to concepts of tax fairness, but I think it applies similarly to concepts of tax efficiency. My first example is the idea of a Pareto-improvement, which unavoidably relies on a comparative baseline. Where the theory and practice of Pareto-optimization in taxation takes a status quo distribution as the baseline for comparison, it also exhibits a baseline problem. My second example is the idea of deadweight loss, which suggests that the posttax world reflects lost opportunities to achieve mutual benefit. It implies that failure to capture the whole market surplus is a cost society must bear in order to have taxation. Yet the notion of the whole market surplus makes little sense in practical contexts, since it assumes that maximum gain is achieved with no taxation at all. This is clearly false, as such a world lacks the legal infrastructure necessary for modern market activity, which is what makes significant market surplus possible in the first instance.

Problematic baseline thinking is not only evident in theory, but also in the practice of tax policy analysis and the public discourse about tax justice. A paradigm case of the former is the now international practice of reporting tax expenditures. This practice relies on measuring nominal revenues foregone when the real tax system is compared with a hypothetical alternative. In order to appear neutral and principled, that hypothetical is construed to be as value-free and

objective as possible. Yet in the end this only masks its obvious normative content. A more neutral and principled approach would candidly articulate and defend the most influential normative considerations that animate different conceptions of the tax system. It would compare the real tax system to multiple ideals, revealing how different principles result in entirely different expenditure accounting. Such accounting is explicitly principled, and it is more neutral than the alternative which simply selects one benchmark and then obfuscates its normative commitments.

Problematic baseline thinking is also evident in public discourse, as nearly everyone conceives of taxation as reducing their pretax entitlements. The primary normative question is, as the public finance orthodoxy articulates, how to spread the cost in a fair and efficient way. In the public mind the emphasis is on everyone paying a fair share. On my view, this makes little sense, since the baseline from which taxes are “paid” does not represent an accounting of entitlements. Moreover, what matters is the way the social system as a whole impacts our lives. Of course, taxation is an important part of that consideration, since it is a major element in determining the distribution of economic power. The notion of a fair share is therefore relevant after all, but the important question for tax justice is whether people receive a fair share, rather than whether they pay it.

This way of thinking about tax justice is, in some ways, disappointing. It is mostly critical—a diagnosis of how *not* to think about taxation. It does not fully answer Piketty’s initial call for a clear statement of alternatives. Moreover, it rejects as myopic the very idea of tax justice as a distinct domain. It implies that tax justice is a kind of misnomer. There is no justice in taxation without distributive justice more generally. As tax justice is subsumed by that larger project, we face all the difficult and enduring questions about who is entitled to what, and why.

Unfortunately questions of distributive justice cannot be partitioned off and resolved in the domain of taxation independently. Yet Piketty also emphasizes that the concepts and frameworks we adopt in addressing these questions can make all the difference. The way we think about property, ownership, efficiency, and fairness will dictate our conclusions about the justice of our overall economic system, including taxation. Unimaginable as it may now seem, I think we would be well-served to abandon the psychologically powerful yet morally bankrupt idea that preliminary descriptions of juridical entitlements track something profound about the way we as a collective ought to relate to each other and to the resources with which we build our lives.

## References

- Acheson, J. (1988). *The Lobster Gangs of Maine*. University Press of New England.
- Ackerman, B (1977). *Private Property and the Constitution*. Yale University Press.
- Alchian, A. (1977). *Economic Forces at Work*. Liberty Fund Inc.
- Anderson, E. (1999). What Is the Point of Equality? *Ethics*, 109(2), 287–337.
- Arendt, H. (1951). *The Origins of Totalitarianism*. Meridian Books.
- Attas, D. (2005). Fragmenting Property. *Law and Philosophy*, 25(1), 119-149.
- Bailey, M. (1992). Approximate Optimality of Aboriginal Property Rights. *The Journal of Law and Economics*, 35(1), 183-198.
- Benson, P. (2002). Philosophy of Property Law. In Coleman, J., & Shapiro, S. (Eds.), *The Oxford Handbook of Jurisprudence and Philosophy of Law* (752-814). Oxford University Press.
- Bentham, J. (2012). *Theory of Legislation (Classic Reprint)*. Forgotten Books.
- Bentham, J. (1864). *Principle of the Civil Code*.
- Biron, L. (2018). Taxing or Taking? Property Rhetoric and the Justice of Taxation. In O’Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (81-97). Oxford University Press.
- Blackstone, W. (1776). *Commentaries on the Laws of England*.
- Bowles, S. (2006). *Microeconomics: Behavior, Institutions, and Evolution*. Princeton University Press.
- Boyce, B. (2007). Property as a Natural Right and as a Conventional Right in Constitutional Law. *Loyola of Los Angeles International and Comparative Law Review*, 29(2), 201-290.
- Bradford, D. F. (1986). *Untangling the Income Tax*. Harvard University Press.

- Bradford, D. F. (1995). *Distributional Analysis of Tax Policy*. AEI Press.
- Breakey, H. (2014). Who's Afraid of Property Rights? Rights as Core Concepts, Coherent, Prima Facie, Situated, and Specified. *Law and Philosophy*, 33, 573-603.
- Brenkert, G. (1998). Self-Ownership, Freedom, and Autonomy. *The Journal of Ethics*, 2(1), 27-55.
- Brennan, G. (1984). Tax Reform and Tax Limits: Political Process in Public Finance. *Australian Tax Forum*, 1(1), 83-95.
- Brennan, G. & Buchanan, J. (2000). *The Power to Tax: Analytical Foundations of a Fiscal Constitution*. Liberty Fund.
- Brennan, G. (2005a). "The Myth of Ownership" Liam Murphy and Thomas Nagel. *Constitutional Political Economy*, 16(2), 207–219.
- Brennan, G. (2005b). The Myth of Ownership. *Australian Journal of Legal Philosophy*, 30, 129-132.
- Brennan, G. (2008). Market Failure: Compared to What? *Ethics and Economics*, 6(1), 1-6.
- Brennan, G. (2018). Striving for the Middle Ground: Taxation, Justice, and the Status of Private Rights. In O'Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (60-80). Oxford University Press.
- Brooks, N. (1986). "Review of Tax Expenditures, by S. S. Surrey and P. R. McDaniel." *Canadian Tax Journal*, 34, 681-691.
- Buchanan, J. (1976). Taxation in Fiscal Exchange. *Journal of Public Economics*, 6(1-2), 17-29.
- Buckle, S. (1991). *Natural Law and the Theory of Property: Grotius to Hume*. Oxford University Press.
- Cairns, H. (1935). *Law and the Social Sciences*. Harcourt, Brace & Company.

- Christman, J. (1991). Self-Ownership, Equality, and the Structure of Property Rights. *Political Theory*, 19(1), 28-46.
- Claeys, E. (2011). Bundle-of-sticks Notions in Legal and Economic Scholarship. *Econ Journal Watch*, 8(3), 205-214.
- Cohen, F. (1935). Transcendental Nonsense and the Functional Approach. *Columbia Law Review*, 35(6), 809-949.
- Cohen, G. (1995). *Self-Ownership, Freedom, and Equality*. Cambridge University Press.
- Compañía General de Tabacos de Filipinas v. Collector of Internal Revenue* [1927] USSC 178; 275 U.S. 87; 48 S.Ct. 100; 72 L.Ed. 177; No. 42 (21 November 1927)
- Cooter, R. (1989). The Best Right Laws: Value Foundations of the Economic Analysis of Law. *Notre Dame Law Review*, 64, 817-837.
- Cordelli, C. (2020). *The Privatized State*. Princeton, NJ: Princeton University Press.
- de Jasay, A. (1997). *Against Politics: On Government, Anarchy, and Order*. Routledge.
- Dean, S. (2009). More Cooperation, Less Uniformity: Tax Deharmonization and the Future of the International Tax Regime. *Tulane Law Review*, 84, 125-166.
- Demsetz, H. (1967). Toward a Theory of Property Rights. *The American Economic Review*, 57(2), 347-359.
- Devereux, M. P., & Vella, J. (2014). Are We Heading towards a Corporate Tax System Fit for the 21st Century? *Fiscal Studies*, 35(4), 449–475.
- Dietsch, P. (2015). *Catching Capital: The Ethics of Tax Competition* (1st ed.). Oxford University Press.
- Dietsch, P., & Rixen, T. (2019). Debate: In Defence of Fiscal Autonomy: A Reply to Risse and Meyer. *Journal of Political Philosophy*, 27(4), 499–511.

- Dorfman, A., & Harel, A. (2016). Against Privatization as Such. *Oxford Journal of Legal Studies*, 36, 400-427.
- Duff, D., Alarie, B., Loomer, G., & Phillips, L. (2018). *Canadian Income Tax Law*, 6<sup>th</sup> ed. LexisNexis Canada.
- Dworkin, R. (2002). *Sovereign Virtue: The Theory and Practice of Equality*. Harvard University Press.
- Econometrics Laboratory at Berkeley (undated). Basic Welfare Economics and Optimal Tax Theory.
- Einhorn, R. L. (2006). *American Taxation, American Slavery*. University of Chicago Press.
- Elkins, D. (2006). Horizontal Equity as a Principle of Tax Theory. *Yale Law & Policy Review*, 24(1), 43-90.
- Ellickson, R. (1986). Of Coase and Cattle: Dispute Resolution Among Neighbors in Shasta County. *Stanford Law Review*, 38(3), 623-687.
- Ellickson, R. (1994). *Order without Law*. Harvard University Press.
- Ellickson, R. (2013). Stone-Age Property in Domestic Animals: An Essay for Jim Krier. *Property Rights Conference Journal* 2(001), 1-25.
- Epstein, R. (1986). Taxation in a Lockean World. *Social Philosophy and Policy*, 4(1), 49-74.
- Epstein, R. (2011). Bundle-of-Rights Theory as a Bulwark Against Statist Conceptions of Private Property. *Econ Journal Watch* 8(3): 223-235.
- Estlund, D. (2019). *Utopophobia*. Amsterdam University Press.
- Frajman, E., & Oppenheimer, J. (2003). The Myth of Ownership: Taxes and Justice. *Economics and Philosophy*, 19, 345-387.
- Feldstein, M. (1976). On the Theory of Tax Reform. *Journal of Public Economics* (6), 77-104.

- Fisher, A. & McClennen, E. (2011). The Pareto Argument for Inequality Revisited.
- Fisher, I. (2007). *The Nature of Capital and Income*. Cosimo Classics.
- Fleurbaey, M. (2018). Welfarism, Libertarianism, and Fairness in the Economic Approach to Taxation. In O'Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (37-59). Oxford University Press.
- Fowler, T. (2015). Marks, Choice and Agency. *Res Publica*, 21, 347-361.
- Freyfogel, E. (2007). *On Private Property*. Beacon Press.
- Freeman, S. (2011). Capitalism in the Classical and High Liberal Traditions. *Social Philosophy and Policy*, 28(2), 19-55.
- Fried, B. (1995). Wilt Chamberlain Revisited: Nozick's "Justice in Transfer" and the Problem of Market-Based Distributions. *Philosophy & Public Affairs*, 24(3), 226-245.
- Fried, B. (1999). The Puzzling Case for Proportionate Taxation. *Chapman Law Review*, 2(1), 157-195.
- Fried, B. (2004). Left-Libertarianism: A Review Essay. *Philosophy and Public Affairs*, 32(1), 66-92.
- Fried, B. (2005). Left-Libertarianism, Once More: A Rejoinder to Vallentyne, Steiner, and Otsuka. *Philosophy and Public Affairs*, 33(2), 216-222.
- Fried, B. (2018). The Case for a Progressive Benefits Tax. In O'Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (147-166). Oxford University Press.
- Friedman, M. (1962). *Capitalism and Freedom*. University of Chicago Press.
- Gallie, W.B. (1956). Essentially Contested Concepts. *Proceedings of the Aristotelian Society*, 56, 167-198.

- Gaus, G. (2010). Coercion, Ownership, and the Redistributive State: Justificatory Liberalism's Classical Tilt. *Social Philosophy and Policy*, 27(1), 233-275.
- Gaus, G. (2012). Property. In Estlund, D. (Ed.), *The Oxford Handbook of Political Philosophy* (93-112). Oxford University Press.
- Gibbard, A. (1976). Natural Property Rights. *Noûs*, 10(1), 77-76.
- Glackin, S. (2014). Back to Bundles: Deflating Property Rights, Again. *Legal Theory*, 20, 1-24.
- Goodman, N. (1976). *Reconceptions in Philosophy and Other Arts and Sciences*. Routledge.
- Graeber, D. (2011). *Debt: the First 5,000 Years*. Melville House Publishing.
- Gretton, G. (2007). Ownership and its Objects. *The Rabel Journal of Comparative and International Private Law*, 71, 802-851.
- Grey, T. C. (1980). The Disintegration of Property, *American Society for Political and Legal Philosophy*, 22, 69-85
- Griffin, J., (1986) *Well-Being: Its Meaning, Measurement, and Moral Importance*. Oxford: Oxford University Press.
- Habermas, J. (1979). *Communication and the Evolution of Society*. Beacon Press.
- Hacker, J. (2002). *The Divided Welfare State: The Battle over Public and Private Social Benefits in the United States*. Cambridge University Press.
- Haig, R. M. (1921). *The Federal Income Tax*. Columbia University Press.
- Halliday, D. (2013). Justice and Taxation. *Philosophy Compass*, 8(12), 1111-1122.
- Hamlin, A. (2018). What Political Philosophy Should Learn from Economics. In O'Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (17-36). Oxford University Press.
- Harberger, A. (1971). Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay, *Journal of Economic Literature*, 9(3), 785-797.

- Hardin, G. (1968). The Tragedy of the Commons. *New Series*, 162(3859), 1243-1248.
- Hausman, D., McPherson, M., & Satz, D. (2016). *Economic Analysis, Moral Philosophy, and Public Policy*. Cambridge University Press.
- Hayek, F. A. (1960). *The Constitution of Liberty*. University of Chicago Press.
- Hayek, F. A. (2014). *The Road to Serfdom: Text and Documents - the Definitive Edition*. Routledge.
- Head, J. (1993). Tax-Fairness Principles: A Conceptual, Historical, and Practical Review. In Maslove, A. (Ed.), *Fairness in Taxation: Exploring the Principles*. (3-55). University of Toronto Press.
- Heaman, H., & Tough, D. (Eds.) (2020). *Who Pays for Canada? Taxes and Fairness*. McGill-Queen's University Press.
- Heath, J. (2020). *The Machinery of Government: Public Administration and the Liberal State*. Oxford University Press.
- Hellmuth, W., & Oldman, O. (Eds.) (1973). *Tax Policy and Tax Reform: Selected Speeches and Testimony of Stanley S. Surrey*. Commerce Clearing House.
- Hobbes, T. (1651). *De Cive*.
- Hockett, R., & James, A. (2020). *Money From Nothing: Or, Why We Should Stop Worrying About Debt and Learn to Love the Federal Reserve*. Melville House.
- Hoebel, E. A. (1954). *The Law of Primitive Man*. Harvard University Press.
- Hohfeld, W. N. (1917). Fundamental Legal Conceptions as Applied in Judicial Reasoning. *The Yale Law Journal*, 26(8), 710-770.
- Hollis, M. & Smith, S. (1994). Two Stories about Structure and Agency, *Review of International Studies*, 20(3), 241-251.

- Holmes, O. W., (2011). *The Common Law*. University of Toronto Law School Typographical Society.
- Holmes, S., & Sunstein, C. R. (2000). *The Cost of Rights: Why Liberty Depends on Taxes*. W. W. Norton & Company.
- Honoré, A. M. (1961). Ownership. In Smith, P. (Ed.), *The Nature and Process of Law: An Introduction to Legal Philosophy* (370-375). Oxford University Press.
- Howard, C. (2009). Making Taxes the Life of the Party. In Martin, I., Mehrotra, A., & Prasad, M. (Eds.), *The New Fiscal Sociology: Taxation in Comparative and Historical Perspective* (86-100). Cambridge University Press.
- Howard, C. (1997). *The Hidden Welfare State: Tax Expenditures and Social Policy in the United States*. Princeton University Press.
- Hume, D. (1739). *A Treatise of human Nature*. Project Gutenberg eBook.
- Hussain, W. (2018). The Common Good. *The Stanford Encyclopedia of Philosophy*, (Spring 2018 Edition), Edward N. Zalta (ed.).
- Hutchison, E. (2017). Principles of Microeconomics.  
<https://pressbooks.bccampus.ca/uvicecon103/>
- James, E. J., & Taussig, F. W. (1886). The State as an Economic Factor. *Science*, ns-7(173S), 485-491.
- James, W. (1907). *Pragmatism: A New Name for Some Old Ways of Thinking*. Longman Green and Co.
- Jamieson, D. (1975). David Lewis on Convention. *Canadian Journal of Philosophy*, 5(1), 73-81.
- Jones, J. W. (1929). Expropriation in Roman Law. *Law Quarterly Review*, 45(4), 512-527.

- Kane, M. A. (2015). A Defense of Source Rules in International Taxation. *Yale Journal on Regulation*, 32(2), 311-361.
- Kant, I. (1797). *Metaphysical Elements of Justice*.
- Kant, I. (1797). *Metaphysics of Morals*.
- Kaplow, L. (1989). Horizontal Equity: Measures in Search of a Principle. *National Tax Journal*, 42(2), 139-154.
- Kaplow, L. (1994). Human Capital Under an Ideal Income Tax. *Virginia Law Review*, 80(7), 1477-1514.
- Kaygusuz, R. (2010). Taxes and Female Labour Supply. *Review of Economic Dynamics*, 13(4), 725-741.
- Kelton, S. (2020). *The Deficit Myth: Modern Monetary Theory and the Birth of the People's Economy*. Public Affairs.
- Kekes, J. (2010). The Right to Private Property: A Justification. *Social Philosophy and Policy*, 27(1), 1-20.
- Kern, A. (2020). Illusions of Justice in International Taxation. *Philosophy & Public Affairs*, 48(2), 151-184.
- King, M. (1983). An Index of Inequality: With Applications to Horizontal Equity and Social Mobility. *Econometrica*, 51(1), 99-115.
- Klein, D., & Robinson, J. (2011). Property: A Bundle of Rights? Prologue to the Property Symposium. *Econ Journal Watch* 8(3): 193-204.
- Krier, J.E. (2009). Evolutionary Theory and the Origin of Property Rights. *Cornell Law Review*, 95, 139-159.

- Kronman, A. (2017). *Education's End: Why Our Colleges and Universities Have Given Up on the Meaning of Life*. Yale University Press.
- Lametti, D. (1998). Property and (Perhaps) Justice. *McGill Law Journal*, 43, 663 – 727.
- Lester, J. (2019). Finances of the Nation: Tax Expenditures in Canada—Historical Estimates and Analysis. *Canadian Tax Journal*, 67(3), 755-773.
- Lewis, D. (1969). *Convention*. Wiley-Blackwell.
- Lewis, D. (1986). *On the Plurality of Worlds*. Wiley-Blackwell.
- Lipsey, R.G., & Lancaster, K. (1956). The General Theory of the Second Best. *The Review of Economic Studies*, 24(1), 11-32.
- Locke, J. (1960). “The Second Treatise of Government.” *Two Treatises of Government*. Edited by Peter Laslett. Cambridge University Press.
- Locke, J. (1980). *Second Treatise of Government*. Hackett Publishing Company, Inc.
- Lockwood, B. B., & Taubinsky, D. (2017). Regressive Sin Taxes. *National Bureau of Economic Research Working Paper Series*.
- Maclean, N. (2017). *Democracy in Chains*. Viking Press.
- Macleod, C. M. (2008). Market Failure, Justice, and Preferences. *Ethics and Economics*, 6(1), 1-7.
- Mandeville, B. (2022). The Harmonization of Tax Legislation: Asymmetry Versus Uniformity Part I. Department of Justice Canada.
- Mankiw, N., Weinzierl, M., & Yagan, D. 2009. Optimal Taxation in Theory and Practice. *Journal of Economic Perspectives*, 23(4): 147-174.
- McDaniel, P., & Surrey, S. (1985). *Tax Expenditures*. Harvard University Press.
- McGuire, R. (2003). *To Form a More Perfect Union*. Oxford University Press.

- McKechnie, W. S. (1896). *The State & the Individual: An Introduction to Political Science, with Special Reference to Socialistic and Individualistic Theories*. Macmillan & Co.
- Merrill, T., & Smith, H. (2001). What Happened to Property in Law and Economics? *The Yale Law Journal*, 111(2), 357-398.
- Merrill, T., & Smith, H. (2007). The Morality of Property. *William & Mary Law Review*, 48, 1849-1895.
- Merrill, T. (2012). The Property Strategy. *University of Pennsylvania Law Review*, 160, 2061-2095.
- Mill, J. S. (1909). *Principles of Political Economy with Some of their Applications to Social Philosophy* (W.J. Ashley, Ed.). Longmans, Green, and Co.
- Mill, J. S. (1979). *Socialism*. Belfords, Clarke, & Company.
- Mirrlees, J. (1971). An Exploration in the Theory of Optimum Income Taxation. *The Review of Economic Studies*, 38(2), 175-208.
- Mueller, D. C. (2003). *Public Choice III* (3rd ed.). Cambridge University Press.
- Murphy, L., & Nagel, T. (2002). *The Myth of Ownership: Taxes and Justice*. Oxford University Press.
- Murphy, L. (2005). Author's Introduction and Author's Response to the Commentators. *Australian Journal of Legal Philosophy*, 30, 122-128; 147-154.
- Musgrave, R. (1959). *The Theory of Public Finance: A Study in Public Economy*. Tata-McGraw-Hill Book Company.
- Musgrave, R., & Musgrave, P. (1989). *Public Finance in Theory and Practice* (fifth edition). McGraw-Hill book Company.
- Musgrave, R. (1990). Horizontal Equity, Once More. *National Tax Journal*, 43(2), 113-122.

- Narveson, J. (1998). Libertarianism vs. Marxism: Reflections on G.A. Cohen's *Self-Ownership, Freedom and Equality*. *The Journal of Ethics*, 2: 1-26.
- Narveson, J. (2010). Property and Rights. *Social Philosophy and Policy*, 27(1), 101-134.
- Nicholas, B. (2008). *An Introduction to Roman Law*. Oxford University Press.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Basic Books.
- O'Donnell, T. (2011). The Restitution of Holocaust Looted Art and Transitional Justice: The Perfect Storm or the Raft of the Medusa? *The European Journal of International Law*, 22(1), 49-80.
- Olson, K. A. (2010). The Endowment Tax Puzzle. *Philosophy & Public Affairs*, 38(3), 240–271.
- O'Neill, M., & Orr, S. (2018). *Taxation: Philosophical Perspectives*. Oxford University Press.
- O'Neill, M. (2021). Social Justice and Economic Systems: On Rawls, Democratic Socialism, and Alternatives to Capitalism. *Philosophical Topics*, 49(1), 1-35.
- Page, K. (2020). Tax Transparency and Perceptions of Fairness: What It Means, How to Get It, and Why It Matters. In Heaman, H., & Tough, D. (Eds.) *Who Pays for Canada? Taxes and Fairness*. McGill-Queen's University Press.
- Panichas, G. (1983). Hume's Theory of Property. *Philosophy of Law and Social Philosophy*, 69(3), 391-405.
- Panitch, L. (1993). Beyond the Crisis of the Tax State? From Fair Taxation to Structural Reform. In Maslove, A. (Ed.), *Fairness in Taxation: Exploring the Principles*. (135-157). University of Toronto Press.
- Prasad, M. (2011). Tax 'Expenditures' and Welfare States: A Critique. *The Journal of Policy History*, 23(2), 251-266.
- Penner, J. E. (2000). *The Idea of Property in Law*. Oxford University Press.

Petrella, C. (2017). Wealth, Slavery, and the History of American Taxation. *Black Perspectives*.

<https://www.aaihs.org/wealth-slavery-and-the-history-of-american-taxation/>

Pigou, A. C. (1947). *A Study in Public Finance* (Third Edition). Macmillan & Co.

Pigou, A. C. (2016). *The Economics of Welfare*. Wentworth Press.

Piketty, T. (2014). *Capital in the Twenty-First Century*. The Belknap Press of Harvard University Press.

Piketty, T. (2020). *Capital and Ideology*. The Belknap Press of Harvard University Press.

Pipes, R. (1999). *Property and Freedom*. Vintage.

Pozen, D. (2007). Tax Expenditures as Foreign Aid. *Yale Law Journal*, 116(4), 869-880.

Proudhon, P. J. (1840). *What is Property?* Humboldt Publishing Company.

Putnam, H. (2004). *The Collapse of the Fact/Value Dichotomy and Other Essays*. Harvard University Press.

Rawls, J. (1971/1999). *A Theory of Justice*. The Belknap Press of Harvard University Press.

Rawls, J. (1981). *The Basic Liberties and Their Priority*. The Tanner Lectures on Human Values.

Rescorla, M. (2019). Convention. *The Stanford Encyclopedia of Philosophy*.

Ripstein, A. (2004). Authority and Coercion. *Philosophy and Public Affairs*, 32(1), 2-35.

Ripstein, A. (2009). *Force and Freedom: Kant's Legal and Political Philosophy*. Harvard University Press.

Risse, M., & Meyer, M. (2018). The Globalized Myth of Ownership and its Implications for Tax Competition. *SSRN Electronic Journal*.

Rousseau, J. (1755). *Discourse on the Origin of Inequality*.

Rousseau, J. (1762). *The Social Contract*.

- Rüfner, T. (2010). The Roman Concept of Ownership and the Medieval Doctrine of *Dominium Utile*. In Cairns, J. W., & du Plessis, P. J. (Eds.). *The Creation of the Ius Commune: From Casus to Regula*.
- Russell, D. (2018). Self-Ownership as a Form of Ownership. In Schmidtz, D., & Pavel, C. (Eds.), *The Oxford Handbook of Freedom* (21-39). Oxford University Press.
- Sadurski, W. (1989). Some Ruminations on the Harm Principle. *Archives for Philosophy of Law and Social Philosophy*, 75(4), 459-467.
- Saez, E., Slemrod, J., & Giertz, S. H. (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, 50(1), 3–50.
- Salter, J. (2001). Hugo Grotius: Property and Consent. *Political Theory*, 29(4), 537-555.
- Sampford, C., & Wood, D. (1991). Tax, Justice and the Priority of Property. *Poznań Studies in the Philosophy of the Sciences and the Humanities*, 23, 181-208.
- Samuels, W. (1981). The Pareto Principle: Another View. *Analyse & Kritik*, 3, 124-134.
- Sandel, M. J. (2013). *What Money Can't Buy: The Moral Limits of Markets* (Reprint ed.). Farrar, Straus and Giroux.
- Sanders, B. (2016, April 24). *A Soda Tax Would Hurt Philly's Low-Income Families*. Phillymag. <https://www.phillymag.com/citified/2016/04/24/bernie-sanders-soda-tax-op-ed/>
- Satz, D. (2010). *Why Some Things Should Not Be For Sale: The Moral Limits of Markets*. New York: Oxford University Press.
- Schoenblum, J. (2003). Myth of Ownership / Myth of Government. *Virginia Tax Review*, 22, 555-587.
- Schumpeter, J. (1954). *History of Economic Analysis*. Oxford University Press.

- Seligman, E. R. (1908). Progressive Taxation in Theory and Practice. *American Economic Association Quarterly*, 9(4), 3<sup>rd</sup> series, 1-334.
- Sen, A. (2017). *Collective Choice and Social Welfare: An Expanded Edition*. Harvard University Press.
- Shoup, C. (1969). *Public finance*. Aldine Publishing Company.
- Simons, H. C. (1938). *Personal Income Taxation: The Definition of Income as a Problem of Fiscal Policy*. University of Chicago Press.
- Simons, H. C. (1946). Federal Tax Reform. *The University of Chicago Law Review*, 14(1), 20-65.
- Sinnott-Armstrong, W. (2023). Consequentialism. *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition)
- Slemrod, J., & Bakija, J. (2017). *Taxing Ourselves: A Citizen's Guide to the Debate over Taxes* (fifth edition). The MIT Press.
- Smith, A. (2017). *The Wealth of Nations*. CreateSpace Independent Publishing Platform.
- Sparks, J. (1856). *The Works of Benjamin Franklin*. Whittemore & Co.
- Statistics Canada (2021). *Effective tax rates and high-income Canadians, 2019*.  
<https://www150.statcan.gc.ca/n1/daily-quotidien/211112/dq211112c-eng.htm>
- Stewart, M. (2005). Taxes and Justice in Context. *Australian Journal of Legal Philosophy*, 30, 133-142.
- Sugin, L. (2004). Theories of Distributive Justice and Limitations on Taxation: What Rawls Demands from Tax Systems. *Fordham Law Review*, 72(5), 1991-2014.
- Surrey, S. (1970). Tax Incentives as a Device for Implementing Government Policy: Comparison with Direct Government Expenditures. *Harvard Law Review* (83), 705-738.

- Tamosi, J. (2012). *Free Market Fairness*. Princeton University Press.
- Thaler, R. H., & Sunstein, C. R. (2021). *Nudge: The Final Edition* (Revised ed.). Penguin Books.
- Tuerck, D. (1967). Constitutional Asymmetry. *Papers on Non-Market Decision Making*, 2, 27-44.
- Vallentyne, P. (1998). Critical Notice of G. A. Cohen, 'Self-Ownership, Freedom and Equality.' *Canadian Journal of Philosophy*, 28.
- Vallentyne, P., Steiner, H., and Otsuka, M. (2005). 'Why Left-Libertarianism Is Not Incoherent, Indeterminate, or Irrelevant: A Reply to Fried', *Philosophy and Public Affairs* 33(2).
- Vallentyne, P. (2018). Libertarianism and Taxation. In O'Neill, M., & Orr, S. (Eds.), *Taxation: Philosophical Perspectives* (99-110). Oxford University Press.
- van Apeldoorn, L. (2016). BEPS, Tax Sovereignty and Global Justice. *Critical Review of International Social and Political Philosophy*, 21(4), 478-499.
- Vogel, K. (1988). The Justification for Taxation: A Forgotten Question. *American Journal of Jurisprudence*, 33, 19-60.
- Waldron, J. (1990). *The Right to Private Property*. Clarendon Press.
- Werning, I. (2007). Pareto Efficient Income Taxation. *MIT and NBER Public Economics Meetings*.
- Wenar, L. (1997). The Concept of Property and the Takings Clause. *Columbia Law Review*, 97(6), 1923-1946.
- Wicksell, K. (1958). A New Principle of Just Taxation. In Musgrave, R., Peacock, A. (Eds.), *Classic in the Theory of Public Finance* (72-118). International Economic Association Series, London.
- Wheeler, S. (2000). Natural Property Rights as Body Rights. *Noûs*, 14(2), 171-193.

Wolff, J. (2002). *Why Read Marx Today?* Oxford University Press.

Wonnell, C. (2001). Efficiency and Conservatism. *Nebraska Law Review*, 80(4), 643-714.

Young, R. (2005). Everyday Libertarianism, Consequentialism and Income Tax. *Australian Journal of Legal Philosophy*, 30, 143-146.

Zelenak, L. (2003). The Myth of Pretax Income. *Michigan Law Review*, 101(6), 2261-2274.