Computer Aided Tax Avoidance Policy Analysis

by

Jacob Rosen

B.S. Mathematics and Complex Economic Systems
University of Michigan, 2012

Submitted to the Engineering Systems Division in partial fulfillment of the requirements for the degree of

Master of Science in Technology and Policy
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Author . . . .

Engineering Systems Division
8 May 2015

Certified by . . . .

Una-May O'Reilly
Principal Research Scientist
Thesis Supervisor

Certified by . . . .

Erik Hemberg
Postdoctoral Associate
Thesis Supervisor

Accepted by . . . .

Dava J. Newman
Professor of Aeronautics and Astronautics and Engineering Systems
Director, Technology and Policy Program

Computer Aided Tax Avoidance Policy Analysis

by

Jacob Rosen

Submitted to the Engineering Systems Division
on 8 May 2015, in partial fulfillment of the
requirements for the degree of
Master of Science in Technology and Policy

Abstract

This thesis presents a three part methodology for analyzing the flow of taxable income in large partnership structures. The method forms the basis for prototypical software which would clarify many complicated basis adjustment issues associated with partnership taxation.

Partnerships, the most common form of "flow-through" tax entities, have rapidly increased in size, complexity and economic relevance between 2005 to 2015, as well as resulting in an estimated $91 billion in underreported income. Many of these partnerships have upwards of one million direct and indirect partners, as well as 100 tiers of additional large partnerships. This surge in the number of partnerships, combined with the highly complicated nature of US partnership taxation law, requires novel techniques to evaluate the tax consequences of increasingly complex financial activity.

A computational methodology is presented in this thesis for understanding and analyzing the allocation of taxable income in large partnership structures, with particular focus on characterizing abusive tax behavior. First, a formal notation is established to fully describe how taxable income is allocated in partnerships, forming the basis of a functioning partnership tax calculator. Next, a simulation is described that processes transaction sequences through partnership structures, as well as a method for assigning audit likelihood to potentially suspicious combinations of financial activity. Finally, a means by which to optimize *a)* transaction sequences that minimize both tax liability and audit likelihood and *b)* auditing procedures that characterize abusive tax behavior in a compact form is established.

The proposed methodology offers taxpayers, auditors and policy-makers a computational approach to resolve uncertainty in partnership taxation, lower the cost of the auditing process through automation and provide a conceptual exploration of tax policy implications.

Thesis Supervisor: Una-May O'Reilly
Title: Principal Research Scientist

Thesis Supervisor: Erik Hemberg
Title: Postdoctoral Associate

Acknowledgments

Working on this project for the past year and a half has been the most intellectually satisfying endeavor of my life. This is a function of both the subject matter and the people I've been surrounded with since September of 2013. Please bear with me as I gush and pontificate.

First and foremost, I need to thank Dr.Erik Hemberg. We spent endless hours in meetings trying to make sense of arguably the most complicated aspect of US tax law. My insight into a computational approach to partnership taxation was honed by his constant inquiry, guidance and micromanagement. Somewhere between a boss, partner and strange older cousin, this thesis is as much his as it is mine.

Dr.Una-May O'Reilly had the brilliant intuition three years ago to model the tax evader-IRS auditor relationship as a co-evolutionary process, a heuristic which continues to influence this project. She has granted me so much more trust and support than I ever could have imagined, let me form my curriculum based on what I believed was important, and guided the project with her deep insight time and time again.

As much as I'd love to pretend that it didn't, this project existed for a year and a half before I hopped on. Sanith Wijesinghe and Geoff Warner, our contacts at MITRE, granted MIT with an impressive amount of freedom from the very beginning to explore the realm of abusive tax shelters, for which I am forever grateful. In the same vein, much of the data structures used in the software implementation were inherited from Osama Badar, the previous version of me on the project.

US tax law and policy is a difficult concept to wrap one's mind around, and it took a lot of instruction until I felt as though I wasn't constantly drowning in a sea of legalese. Professor Stephen Shay has provided invaluable insight into the tax policy world, as well as listening to me gab about my worries conerning my uncertain future. An enormous thank you is owed to Professor Howard Abrams in San Diego for clarifying countless partnership taxation issues for me. He took time out of his days to speak with a random graduate student, and might be the only person on earth who is both interested in the mathematics and can verify whether they are correct. Finally, Professor Daniel Halperin taught my introduction Taxation course in a manner which was not only comprehensible by a non-law student such as myself, but provided a subtle ethical paradigm through which to view the purpose of taxes in modern society.

To the TPP administrators, thank you so much for not murdering me when I refused to take any other research project. It would have been very difficult to write this thesis if you had. I cannot think of any other program that would have allowed me to wait this long to find a consistent source of funding, and I will always remember the amazing opportunities that has afforded me.

For the past two years, I've lived with and studied aside a group of TPP students who have consistently forced me to reevaluate my beliefs and think critically, rather than hanging out in my own comfort zone. I'll always remember late nights with Alex, Chris and Justin [1], drinking beer and arguing about everything from Bashar Al-Assad to the weather in Texas. To all other TPP students in my class, I'm amazed by all of you, and I'm so lucky to have met such a weird, smart, kind and goofy group of nerds.

Chloe, our long talks about gender and sexuality, liberal politics and that other stuff that I shouldn't put in my thesis likely kept me sane when I was staring at computer code for ten hours a day. Thanks for everything.

Grandma and grandpa, I still think about you whenever I make any decision about what to study or pursue in my life. The two of you have provided me with the best moral compass that anyone could ask for, now it's up to me to figure out how it works.

Mom, dad and Rebecca. You're all completely insane in the same way that I am and I love you for it. Without your undying support for me as a human, I never would have had the courage to do any of the things I've accomplished. So you guys get the final thank you, which is arguably the most important.

[1] and John, but he came later

Contents

List of Figures

List of Tables

1

Introduction

The ability that many Americans have to minimize their tax liability by utilizing potentially abusive financial mechanisms not only deprives the United States government of revenue, but perpetuates the impression that citizens' access to justice depends heavily on the amount of resources they have available.

This thesis presents a computational methodology for analyzing and exploring areas of taxation, specifically those concerning partnerships, that may contain inefficiencies or encourage abuse. The proposed approach to policy analysis can identify aspects of the tax code which cause confusion amongst taxpayers, incentivise intentionally abuse behavior, or both.

1.1 Preamble

The US tax code has become extraordinarily complex over the past several decades [10], resulting in increasing levels of inequality when navigating the tax code. Aid to low-income Americans is often granted through the tax code, such as the Earned Income Credit (EIC), which essentially subsidizes workers' wages by offering a refundable credit proportional to the amount that they work. While such programs are successful, they often involve navigating through complicated tax documents that requires trained tax preparers. Additionally, because most low-income tax breaks merely offset taxes paid on their income, there is a low ceiling at which many Americans can take advantage of government programs.[1]

In the same way, government benefits in the form of tax deductions or credits have a much greater effect on the wealthy. Not only does a larger amount of income increase the incentive to mitigate taxes on that income, but there are more resources available to navigate the tax code in search of tax minimizing techniques.

[1] I do not attempt to disentangle the relationship between Americans and taxes, nor do I speak of the disadvantages that tax complexity has on the less economically fortunate. These are worthwhile endeavors that have, and should be explored in a different context.

The complexity of the tax code results in a myriad of unintended consequences arising from interactions between clauses made without regard to one another and often times in different decades. For example, the passive activity loss clause outlined in §469 of the Internal Revenue Code (IRC) was drafted largely in response to leasing schemes that took advantage of such consequences [36]. Generally, these schemes involved wealthy taxpayers purchasing depreciable property from businesses and immediately leasing it back to them for a slightly lower price. The wealthy taxpayers reduced their tax liability, while the business was buying property back for less than they sold it. A canonical example illustrates the mechanics.

Example 1. An airline company Dewey Financial (DF) owns a fleet of ten airplanes, each worth $1 million. The tax code allows owners of physical property such as airplanes to deduct a certain amount per airplane, say $100,000, from their taxable income to account for normal wear and tear. Supposing that the corporate tax rate is 35%, the depreciation deduction from each airplane reduces the company's tax liability by $35,000 by eliminating $100,000 of taxable income.

The problem is that DF has a limit on the amount of deductions it can take, and suppose that they are only allowed to take depreciation deductions on five of their ten planes. While the company itself cannot the deductions, there is an opportunity to essentially sell the deductions to high-income taxpayers. And given that the tax rate on the richest taxpayers was 70%, a $100,000 deduction can save them $70,000.

DF thus finds a taxpayer Terry with $10 million in taxable income, who agrees to purchase five airplanes for $1,017,500 each. Terry then leases back the planes as if they each cost $982,500. Once Terry takes the depreciation deductions on the planes, both she and DF have made a $175,000 profit, funded by the US government. [2]

The depreciation deductions described in the above example introduces the concept of a "passive loss", which is a deduction taken from an activity "in which the taxpayer does not materially participate" (§469$(c)(1)(B)$). Deductions resulting from passive losses were disallowed by the 1986 Tax Reform Act, but it illustrates how different clauses, in this case depreciation deductions and definition of Adjusted Gross Income (AGI), can interact to produce perverse results. But given the number of laws governing taxation, there exist plenty of other unintended consequences that policy-makers would never have considered. This presents an opportunity for those with a deep understanding of the tax law to exploit underlying flaws in the Internal Revenue Code for their clients' benefit.

An anecdote clarifying manipulation of the tax code involves the childhood "no backsies"

[2] DF sells all five planes for a combined $5,087,500 and buys them back through a lease for $4,912,500, resulting in a $5,087,500−$4,912,500 = $175,000 profit. While Terry at first loses $175,000 from the leasing agreement, she is able to claim a total of $500,000 in depreciation deductions, which at a 70% tax rate, results in $350,000 in tax savings. Thus, her total profit from the endeavor is $350,000 − $175,000 = $175,000 as well.

rule [3]. The rule stipulates that if there is a group of children in a line, then one can allow friend to enter the line in front of, but not behind them. The moral justification for the rule is that if everyone in the line suffers from the extra wait time, then the child that let their friend cut in line should suffer as well. But this rule is easily evaded if, immediately upon letting their friend cut in front of them, the child exits the line. In turn, the child's friend allows them to legally cut in line in front of them, effectively engineering a "backsie" from two legal actions.

In the same vein, by encoding aspects of the tax law into a computer program, the IRS can iterate through possible combinations of tax rules to determine their often times unknown tax consequences. Given that many of the laws were written without regard to one another and often times in difference decades, they have the capacity to interact adversely. Thus, creating a tool to explore the cause and effect of both pre-existing and potential tax laws can provide significant value.

This thesis will present a method by which legal professionals can evaluate partnership tax law to explore transactions implemented through partnership structures with a pure tax motive. Some of these transactions are perfectly legal, others are illegal, and many exist in a gray area in between. The goal is to shed light on potentially unintended consequences that are embedded deep within the Code so that they can be rectified in order to prevent future iterations of abusive tax behavior.

1.2 Motivations

The subject matter presented and methodology proposed in this thesis resulted from a set of motivating questions.

- How can the US government construct policy in order to disallow current abuses of the tax code?
- Is there a method by which the government can *anticipate* abuses of the tax code before they have the chance to drain billions of dollars in revenue?
- How can the government make preventative, rather than purely reactionary, tax abuse policy?

The tax implications arising from proposed policies, and even many pre-existing stipulations in the Internal Revenue Code, are often times unknown by taxpayers, auditors and policy-makers alike. Laws governing partnership taxation are particularly susceptible to this issue, given that the rules are both heavily dependent on a set of arithmetic calculations that, while individually simple, can interact in a highly complex manner.

[3]Taken from a discussion with Ameek Ponda J.D., LL.M. of Sullivan & Worchester on October 28, 2014

All relevant parties could thus benefit greatly from a computational tool that allows them to evaluate the tax implications arising from various types of common financial activity that is implemented through partnerships. Taxpayers could ensure compliance with the tax law, auditors could better understand financial activities within large partnership structures, and policy-makers could evaluate the effects that proposed law changes could have to the tax system.

1.2.1 Research Questions

The following questions guided the development of the computational tools developed later in this thesis.

1) How can taxable income arising from transactions within a network of partnerships be represented abstractly and subsequently calculated?

2) How can we simulate the process of an IRS auditor, under resource constraints, determining the aspects of a sequence of transactions which indicate suspicious behavior?

3) Can new forms of abusive tax behavior be anticipated by optimizing both tax planning and auditing policy?

   3.1) How can transaction sequences be constructed which minimize tax liability, while attracting minimal IRS suspicion?

   3.2) How can IRS policy be formed which assigns high auditing likelihood to suspicious transaction sequences, without falsely auditing routine transactions.

1.3 Contribution

The methodology proposed in this thesis, and also published in [33, 25, 16], contributes to three separate, but interconnected realms of taxation.

Taxpayers For stakeholders in partnerships, as well as those that prepare their taxes, clarification of the tax consequences resulting from financial activity decreases the amount of resources necessary to ensure compliance with the tax code, as well as likelihood of being audited. Additionally, allowing businesses to evaluate the tax liability incurred from potential transactions could contribute greatly to earning estimations.

Auditors Few auditors at the Internal Revenue Service have the proper training or background to conduct examinations on large, complex partnership returns without additional input from partnership taxation domain experts. A user-friendly piece of software would not only allow auditors to computationally evaluate large partnership structures, but could serve as a cost-effective instructional aid for learning partnership taxation.

Policy-makers Those at both Congress and the IRS, are often times unsure of the effect that a policy change could have on the tax system. Introducing new rules into the proposed software could increase policy-makers' foresight into the operation of the partnership taxation environment. Additionally, many inefficient policies, such as §754 elections and electronic filing requirements, are in place because of the partnership tax return complexity. Such policies could be reconsidered with the introduction of a software standard by which partnerships could file their returns.

1.3.1 Thesis Overview

Chapter 2 will provide background information necessary to understand the proposed methodology. Section 2.1 describes largely mathematical taxation concepts used later in the thesis, focusing specifically on partnership taxation. Section 2.2 will discuss the legality of different types of tax behavior, and attempt to define a lexicon by which these types can be described.

The following four chapters propose and discuss the overarching methodology. First, Chapter 3 will define a mathematical abstraction of partnership taxation, which is necessary to properly calculate taxable income in complex scenarios. Next, Chapter 4 describes a simulation framework by which *(1)* transaction sequences are processed through partnership structures and, *(2)* likelihood of an audit is assigned to the aforementioned financial activity. Chapter 5 then establishes a search space and hypothetical objective function, such that both effective tax evasion schemes and optimal IRS auditing policies can be found with the proper meta-heuristic. The three methods are used to analyze an abusive tax shelter known as Distressed Asset Debt (DAD) in chapter 6.

The potential implications of the proposed methodology are discussed in chapter 7, followed by concluding remarks and potential future work in chapter 8.

2

Background

In order to understand the methodology proposed in Chapters 3, 4 and 5, a brief overview of taxation concepts will be covered. Section 2.1 begins by covering basic calculations regarding taxable income, then eases the reader into complex partnership taxation examples through a series of examples. Next, section 2.2 establishes a lexicon to describe different types of tax behavior, as well as covering several responses that the government has available to prevent the propagation of abusive tax shelters and other illegal activity.

2.1 Taxation Basics

This thesis focuses on taxable income incurred during the sale or trade of property, rather than income received in exchange for services. When an individual receives $\$X$ in wages, the entire $\$X$ is treated as income and taxed as such. Conversely, suppose that same taxpayer were to purchase property, such as a house or a tradable security, for $\$Y$ and later sell it for $\$Z$. Provided $\$Z > \$Y$, then the taxpayer received income of $\$Z - \$Y$, which we refer to as *gain*.

Instead suppose that $\$Y > \$Z$. In this case, the taxpayer actually sold their property at a loss, which should be reflected in their income. Thus, in certain cases, the taxpayer is able to *deduct* $|\$Z - \$Y|$ from their taxable income, thus reducing their tax liability. While the property may have been bought years before it was sold, the sale is reflected on the taxpayer's financial statement the year that it was sold.

Fair Market Value (FMV) The price that an asset or a piece of property is worth on the market, given that all market participants have perfect information and provided that the asset may provide future income streams. The FMV in the example above is $\$Z$.

Adjusted Basis The cost that a taxpayer or entity incurred in acquiring a piece of property after adjusting for other consequential costs. While the word **basis** is used to define the *original* cost, adjusted basis is more accurate for tax purposes. An example of a basis adjustment is when a taxpayer claims depreciation deductions on a piece of property, in which case the adjusted basis is decreased. Basis adjustments will be discussed more in depth later. The adjusted basis in the example above is $\$Y$.

Gain/Loss The difference between a piece of property's FMV and adjusted basis upon sale or disposition[1] of the property. Gain is generally added to the seller's taxable income, and loss can often be subtracted [2]. The gain/loss in the example above is $\$Z - \$Y$.

2.1.1 Partnership Taxation

Partnerships are the most common form of "flow-through" tax entities, which are named as such because they pass any gain or loss incurred during business activity to the partners. This is distinguished from C corporations, which is treated as a taxpayer in the sense that the corporation itself is liable for any taxes, not the individual taxpayers. While this often allows for simpler accounting and can be beneficial to many small business owners, the laws that govern partnership taxation can become very complicated.

A partnership is generally formed when a group of two or more taxpayers, pre-existing partnerships or other entities, decide that they would like to engage in some sort of business activity with one another. The arrangement is made in the *partnership agreement*, which specifies the ownership relations between the partners.

Once the partnership is formed, there are a handful of activities that are important for tax purposes.

1. **Sale of Property:** When a partnership sells property that it owns, there is likely to be either gain or loss arising from sale. In this case, the tax liability is distributed amongst the partners in a manner prescribed by the tax law and the partnership agreement.

2. **Sale of Partnership Interest:** A current partner can choose to exit the partnership or reduce their share in the partnership by selling all or part of their interest to another taxpayer or entity. Tax consequences arise in this case when the FMVs of the assets held by the partnership have changed.

3. **Distribution:** A current partner can also choose to exit or reduce their share in the partnership by receiving a liquidating or non-liquidating distribution. The partner-

[1] This thesis discusses only the *sale* of property

[2] Different types of property often warrant different tax rates, and so the gain or loss resulting from the sale must be added to or subtracted from the correct income pool. The rules determining property types and their associated rates can become very complicated, and will not be discussed in depth in this thesis. For clarification of capital assets for individual taxpayers, see §§1221 and 1222 of the IRC. For similar clarification in respect to property held by partnerships, see §751 of the IRC.

ship transfers some of its property to the partner and the partner's share is reduced. Generally this results in an increase in the remaining partners' shares, and can result in immediate tax consequences.

4. **Contribution:** A taxpayer or entity can become a partner or increase their share by contributing assets or services to the partnership. Generally this results in a decrease in the other partners' shares. While there are no immediate tax consequences, contributions can result in certain important bookkeeping measures.

   Many complicated tax issues in partnerships arise due to the fact that contributions are non-taxable. That is, a partner is not allocated any taxable income upon the contribution of cash or non-cash property. Thus, the adjusted basis in any non-cash property must continue to reflect any pre-contribution gain or loss that the property experienced while owned solely by the contributing partner. This concept is referred to as *carryover basis* because the adjusted basis of contributed property "carries over" to the partnership.

Gain or loss derived from any increase or decrease in the contributed property's FMV after contribution is attributable to all of the partners. But the pre-contribution gain or loss must be attributed solely to the contributing partner. The following example illustrates the concept of gain allocation.

Example 2. Alice, Bob and Cathy form a partnership P with equal distributive shares. Alice contributes a stock X with an adjusted basis of \$15 and an FMV of \$50, while both Bob and Cathy contribute \$50 in cash.

A year later, stock X has increase in value to \$200 and the partners decide to sell it. The adjusted basis of stock X is still \$15 because of the *carryover basis*, so the sale produces $\$200 - \$15 = \$185$ in taxable income that must be allocated amongst the partners.

Intuitively, stock X appreciated \$150 while within the partnership, thus each partner is allocated $\frac{1}{3}$, or \$50 of taxable income. Additionally, Alice remains liable for the \$35 of stock X's pre-contribution gain, which brings her allocated taxable income from the sale of the property to $\$50 + \$35 = \$85$.

Thus, Alice, Bob and Cathy are allocated, respectively, \$85, \$50 and \$50 of taxable income, which totals to the \$185 difference between stock X's FMV and adjusted basis.

Problems arise when a partner that contributed non-cash property exits the partnership before the contributed property is sold. Because an interest in a partnership is a piece of property itself, each partner has an adjusted basis in their interest, which is referred to as the *outside basis* [3]. Upon formation of a partnership, each partner's *outside basis* in their interest is the combined adjusted bases of the property that they contributed. If a partner sells their interest to an outside actor, then gain or loss arising from the sale of the contributed property after the contributing partner exists can result in duplication of either taxable income or deductions.

Example 3. Assume the same initial facts as example 2. Alice's *outside basis* in the partnership is \$15, while both Bob and Cathy have an *outside basis* of \$50. Again, suppose that a year later, the FMV of property X has increased to \$200.

Rather than selling property X, suppose that Alice sells her interest in the partnership to a taxpayer Dan. The combined FMV of the partnership assets is \$300 (\$100 in cash and a \$200 piece of property), thus Alice's share is worth \$100. Upon the sale, she is allocated $\$100 - \$15 = \$85$ of taxable income, which is the difference between the FMV of $\frac{1}{3}$ of P and her *outside basis* in the partnership. Note that this is the *exact* amount of taxable income she was allocated in the previous example.

Suppose that immediately after Alice exits, the partnership decides to sell off property X. The basis of the property is still \$15, so the sale generates \$185 of gain that must be allocated amongst the partners. The question is: how should that \$185 be allocated?

Scenario A

Dan has not experienced *any* increase in wealth derived from property X's increase in FMV, thus none of the gain should be allocated to him. Both Bob and Cathy are liable \$92.50 for of taxable income, which is significantly more than the \$50 they would have been allocated if property X had been sold immediately before Alice sold her share.

Scenario B

The fact that Alice sold her interest in the partnership should *not* affect Bob or Cathy's tax liability. Thus, they are both allocated the \$50 in taxable income as in example 2 and Dan is liable for the remaining \$85 of taxable income, even though he did not benefit from the increase in property X's FMV. This is in fact how the taxable income would be allocated without the enactment of the clauses discussed in the following section 2.1.1.1.

[3] This is distinguished from the concept of a partner's *inside basis*, which refers to the portion of each asset's adjusted basis that is attributable the specific taxpayer. This is a highly complex issue that will be tackled with the representation described in section 3.1.

Neither of the scenarios in example 3 result in sensible tax consequences, and it illustrates how without further specification, gain or loss arising from the same piece of property can be duplicated. When Alice sold her interest in the partnership, she indirectly paid tax on both the pre and post contribution gain from property X. That is, the difference between her initial *outside basis* of $15 and the $50 FMV of the property reflected the pre-contribution gain. Then when she sold her interest for $100, there was an additional $50 in taxable income that represents the property's appreciation while owned by the partnership. In other words, $85 of taxable income derived from property X's increase in FMV had been allocated upon the sale of Alice's partnership interest.

2.1.1.1 The §754 Election

The duplication of taxable income, as shown in example 3, is a well known problem in partnership taxation. In August of 1954, several sections to Subchapter K of the Internal Revenue Code were added in order to rectify the issue. If a partnership "makes" an optional §754 election, then in the event of the sale of a partnership interest or a distribution, then the partnership can adjusted the bases of their assets as prescribed in, respectively, §§743 and 734. Additional specifications are included in §755, as well as in the treasury regulations.

Example 4. Assume the same facts as example 3, except suppose that the partnership made a §754 election upon formation. Once Alice sells her share to Dan, a §743(b) basis adjustment is triggered, which requires that the partnership, in essence, adjust property X's basis to account for the taxable income already allocated to Alice, that was derived from property X's increase in FMV.

Thus, the partnership positively adjusts the basis of property X by the $85 of income that Alice was allocated to $100. Now when the partnership sells property X, only $\$200 - \$100 = \$100$ of gain must be allocated amongst the partners. As one would expect, Bob and Cathy are both allocated $50 from the sale, and Dan is not liable for any of the gain.

Basis adjustments such as this become complicated very quickly. Suppose Alice had contributed many different pieces of property to the partnership at different times. Or suppose that instead of taxpayers, all of the partners in example 4 were partnerships as well [4].

The complexity of the basis adjustments is likely the reason that §754 elections are optional [3]. At the time of the sections' codification in 1954, all of the calculations would have been on paper and required vast amounts of resources to properly compute. Additionally,

[4] Numerous scenarios and clarifications for more complex cases are contained in regulations issued by the US Department of Treasury, namely Treas. Reg. §§1.743 − 1 and 1.755 − 1. A similar, but notably different, type of basis adjustment occurs immediately following a distribution, the specifics of which are contained in §734, and Treas. Reg. §§1.734 − 1 and 1.755 − 1.

partnerships at the time were fairly uncommon and those that existed rarely engaged in complex economic activity. Thus, requiring all partnerships to expend resources to make basis adjustments likely seemed cumbersome and unnecessary.

2.1.2 Contemporary Partnerships

While many of the laws that govern partnership taxation were established at a time when both the frequency and the complexity of partnerships were relatively low, this is not the case anymore. Numerous reports by the Government Accountability Office (GAO) have been issued in 2014 alone that highlight both the rapid increase in the number of partnerships and the problems associated with conducting efficient audits [14, 15, 13]. The number of "large partnerships", defined as partnerships with 100 or more partners and at least $100 million in assets, increased from 2,832 in 2002 to 10,099 in 2011. A majority of these partnerships have well-over 1,000 partners, as shown in figure 2.1 below.

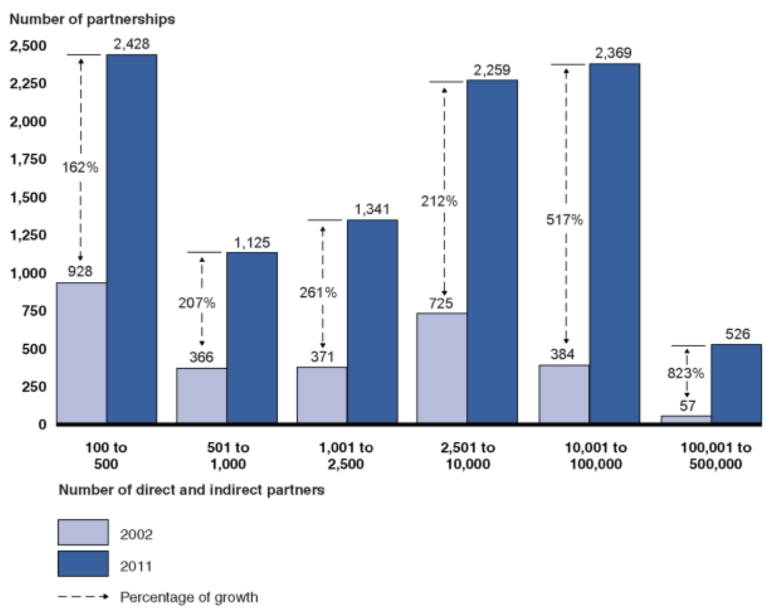

Figure 2.1: Source: GAO analysis of IRS data from Enhanced Large Partnership Indicator File and Business Return Transaction File, Compliance Data Warehouse. GAO-14-732

Contributing significantly to the complexity of large partnerships is the increase in tiered structures. In 2011, 36% of large partnerships had at least 1,000 flow-through entities as direct or indirect partners [14], where an indirect partner is a partner connected through yet another flow-through entity. These types of tiered structures are a large source of confusion for the IRS, 78% of large partnerships had relationships that were six or more tiers deep. Figure 2.2 below illustrates a hypothetical, complex partnership structure.

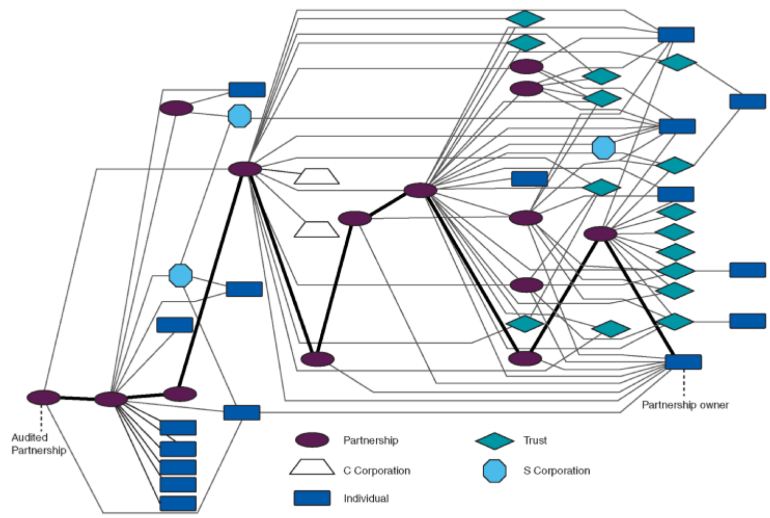

Figure 2.2: Source: GAO analysis of IRS data. GAO-14-732

2.2 Tax Avoidance and Abuse

Necessary for this thesis is the establishment of a lexicon to describe different types of tax minimizing behavior that may or may not violate US tax law. It should be noted that these terms, unless otherwise specified, are not legally recognized. Rather the definitions presented in this section represent the perspective of consulted tax experts [3, 28], as well as convenient distinctions that serve the purpose of framing the methodology proposed in this thesis.

The most natural way to frame the types of potentially abusive activity that this thesis focuses on is to first define what it does *not*.

Tax Evasion A tax reduction plan that, if all of the relevant facts were provided to the US government, would almost certainly result in civil or criminal penalties. Examples range from failing to reporting self-employment income to using foreign banks with the purpose of hiding one's true amount of income [5].

Tax Avoidance A tax reduction plan that, given all of the relevant facts, is either unambiguously legal or highly unlikely to be challenged by the IRS. Prior to the 1986 Tax Reform Act, the airplane leasing arrangement discussed in the introduction classifies as tax avoidance. Other than the fact that policy-makers likely did not intend depreciation deductions to be taken for planes by taxpayers who had no substantive ownership, every aspect of the tax plan, and its result, was perfectly legal.

[5] In this way, the use of the terms "tax evader" or "tax evasion" scheme in chapters 4 and 5 is misleading. While the proposed methodology may lend insight into the determining whether any discrepancies may exist which indicate evasion, it is not the primary focus.

> Tax avoidance thus encapsulates a wide range of behavior. Contributing to charity, using an energy efficient air conditioner and purchasing a home, in their own way, could all be considered tax avoidance by this definition. Making economic decisions based on their tax benefits is a primary driver of US policy and in many cases is intended.

As mentioned in section 1.1, this thesis focuses on tax plans that fall in the gray area between tax evasion and tax avoidance. That is, cases in which, even given all of the relevant facts, are ambiguously legal. Generally such plans are composed of a sequence of steps, that while legal individually, produce a perverse result that may be disallowed. Tax planners are likely aware that the tax outcomes were not intended by law-makers, but have little incentive not to engage in a plan that generates large tax savings with a potentially low probability of being audited.

Consider the example scenarios posed in section 2.1.1 [6].

Example 5. Assume the same facts as example 3, where the partnership has intentionally *not* made a §754 election. Additionally suppose that Alice instead contributes property Y that has an adjusted basis of \$500 and an FMV of \$50. Thus, Alice has an *outside basis* of \$500 in her share of the partnership.

Similarly, suppose that a year later when property Y has increased in value to \$200, Alice decides to sell her interest in the partnership to another taxpayer Emma for \$100. Her *outside basis* in the partnership exceeds the purchase price by \$400, thus she is able to deduct \$400 of income derived from other sources. This is an intended effect, given that she has lost a total of \$400 from her investment in property Y.

Without a §754 election, the adjusted basis of property Y within the partnership remains \$500. When the partnership then decides to sell property Y for \$200, $\$200 - \$500 = -\$300$ of income, or \$300 of deductions, must be allocated amongst the partners.

From a tax perspective, both Bob and Cathy benefited from the \$100 increase in FMV that property Y experienced while held by the partnership, thus, like the previous examples, they are both allocated \$50 of *positive* income. But because a total of $-\$300$ of income must be allocated, this means that Emma claims a \$400 deduction arising from the sale of property.

Thus, in same way that the lack of a §754 election can result in the duplication of taxable gain, it could, prior to 2004, also cause a duplication of deductible loss arising from pre-contribution loss. This is also a perverse result, because Emma did not lose any income

[6] Note that the following example is based on the IRC before the addition of $\S704(c)(1)(C)$ as part of the American Jobs Creation Act of 2004 [32].

as a result in her investment in the partnership or property Y. Results such as this can cause similar situations and deadweight losses as the airplane leasing example posed in the introduction. That is, if Emma has a lot of capital gains elsewhere in her portfolio that she would like to mitigate, Alice might charge her slightly more than the actual price of the share.

On one hand, prior to 2004, this is clearly not tax evasion because every individual step in the example is legal. However, the US government has a set of tools to combat nominally legal activity that they deem abusive. For the purpose of this thesis, will will refer to such tax plans as *abusive tax shelters*, which are dicussed further in the next section.

2.2.1 Tax Shelters

The term "tax shelter", like many of the terms discussed in the previous section, is difficult to universally define. Intuitively, it is a type of tax plan that "shelters" income from the government such that it is not taxed. Many retirement plans such as a 401(k)'s or an IRA classify as a tax shelter, but they exist primarily as a means by which the government incentives certain types of financial activity. For the purpose of this thesis, the term "abusive tax shelter" refers to tax shelters that are not *clearly* intended by the US government [7].

Clearly, this does not offer much guidance towards whether a specific sequence of transactions constitutes a tax shelter. Thus, the government relies on three primary legal statutory and regulatory responses to classify financial activity as a tax shelter [22]. While these responses apply to many types of tax shelters, this thesis will focus on those that apply primarily to plans involving partnerships. The types of responses are described below, in increasing order of generality.

Statutory grants of broad authority

> These are amendments to the IRC, such as clauses in §§269 and 446 which are made in direct response to perceived abusive activity. Generally, such grants are not particularly broad in their application and apply to previously classified behavior [22].

[7] The Internal Revenue Code does include a definition of a tax shelter in §6662(d)(2)(C)(ii) as:

> **(ii)** Tax shelter For purposes of clause (i), the term tax shelter means
>
> **(I)** a partnership or other entity,
>
> **(II)** any investment plan or arrangement, or
>
> **(III)** any other plan or arrangement,
>
> if a significant purpose of such partnership, entity, plan, or arrangement is the avoidance or evasion of Federal income tax.

The amendments to §§704 and 743 in response to the situation described in example 5 illustrate the power of such techniques. In short, they hindered the ability of taxpayers to duplicate deductions by *a)* disallowing the transfer of pre-contribution loss from assets contributed to partnerships and, *b)* requiring a basis adjustment in cases where the sale of a partnership interest would generate a substantially large loss, specifically greater than $\$250,000$. While this measure was effective at preventing the specific type of loss duplication [20], statutory grants of broad authority are generally not forward looking, and ignore the underlying disregard of the intent behind the statutes.

Anti-Abuse Rules

The Department of Treasury released and updates Treas. Reg. $\S 1.701-2$, which outlines activity that may be subject to a court challenge if filed. The document first speaks in very general terms about the application of Subchapter K, then proceeds to present taxpayers with a list of example scenarios that would be subject to the anti-abuse rules. Since the regulations' promulgation in 1994, there has been a notable increase in the US government's ability to disallow certain transactions [29].

Example 9 in the anti-abuse rules attempts to outline abusive scenarios that arise from a partnership's intentional decision to *not* make a §754 election. The first section of the example focuses on a fairly specific case in which particular assets are distributed with the purpose of deferring taxation. But the subsequent sections describe aspects of Subchapter K which imply the necessary making of a §754 election in order to comply with the proper reflection of income. The purpose is to confront scenarios in which the deliberate absence of the election can cause abuse, which has broader applications at the expense of less specificity.

Judicial Doctrines

Over the years, local courts have formed a set of common law doctrines aimed at disallowing clearly abusive tax activity. Many of the doctrines are based off of the **substance over form** doctrine that has its basis in the 1935 Gregory v. Helvering case. The doctrine specifies that a sequence of transactions must demonstrate "economic substance" in order to claim tax benefits. For example, if a taxpayer attempts to treat equity as debt to claim certain deductions, while financial equivalent, the US government may have a right to disallow any such deductions.

Other doctrines are slightly more granular variants on the substance over form doctrine. The step transaction doctrine treats a sequence of steps (transactions) as a single step for tax purposes if it more clearly reflects the economic reality. The business purpose doctrine disallows tax benefits if the financial activity was taken purely for a tax motive. Arguably the most important, the economic substance doctrine is the only codified doctrine modeled after the substance over form doctrine and will be discussed more in depth later in this section.

These tools serve a dual purpose: *1)* to disallow abusive tax shelters as they are realized by the government and *2)* to prevent future abusive tax shelters from propagating. While all of the responses are useful in the combat against abusive activity, the judicial doctrines provide a theoretical foundation for the characteristics of abusive tax behavior.

The most widely used doctrine today is the *economic substance doctrine*, another variant of the substance over form doctrine. Codified as §7701(o) in section 1409 of the Health Care and Education Reconciliation Act of 2010, the doctrine specifies

> In the case of any transaction to which the economic substance doctrine is relevant, such transaction shall be treated as having economic substance only if
>
> **(A)** the transaction changes in a meaningful way (apart from Federal income tax effects) the taxpayers economic position, and
>
> **(B)** the taxpayer has a substantial purpose (apart from Federal income tax effects) for entering into such transaction.

Many have criticized the doctrine for its vague language and lack of clarifying a homogeneous interpretation across courts [27]. But has proven effective at disallowing new iterations of abusive tax shelters [35].

While the focus of this thesis is not the legal techniques used to disallow abusive tax shelters, understanding the methods by which the government combats abusive activity lends insight into the types of behavior that must be detected. That is, it helps to outline the space of tax behavior that lies in the legal gray area between evasion and avoidance as defined in this thesis.

3

Representation

An abstract representation of partnership taxation is necessary to answer the first research question motivating this thesis, namely

How can taxable income arising from transactions within a network of partnerships be represented abstractly and subsequently calculated?

The primary challenge in constructing a mathematical representation of partnership taxation is the conflict between the two pass-thru taxation models:the entity model and aggregate model [1]. In essence, the entity model treats the partnership as a taxable entity, like a corporation, then passes the income or loss through to the partners. While this is sufficient for most partnership activity, there are instances in which a different approach is required. Conversely, the aggregate model does not assign the partnership any tax attributes, but rather treats the partnership as merely an aggregate of its owners. This is more cumbersome and cannot capture many crucial legal aspects of partnerships, but is necessary to track each partners' interest in each individual asset, as shown in the examples in section 2.1.1.

This issue is unique to flow-through entities because §702 specifies that the character of any income, gain, loss, deduction or credit must be realized as if it was "incurred in the same manner as incurred by the partnership". That is, even though the partnership is conducting the financial activity, each partner must treat any income as though they conducted the activity themselves. Corporations avoid this problem because income from business activity is taxed at the entity level, then dividends to each shareholder are treated in respect to the stock they hold, *not* the character of the business activity that generated the income.

Most important for this representation are the provisions described in §751 of the IRC, which distinguish between assets treated as ordinary income vs capital income. For example, income generated from the sale of inventory items of a partnership is treated as ordinary

[1] See appendix C for elaboration from [11] regarding the different models.

for the partners, as opposed to that from a long-term stock, which is capital. The rules become even more complex if the asset being sold was originally contributed by a partner in exchange for an interest. In that case, the pre-contribution gain or loss must be attributed to the contributing partner, while the income generated from any post-contribution change in FMV is allocated amongst all of the partners according to the partnership agreement. This distinction reflects the aggregate model of calculating partnership tax. Unlike corporations, the *character* of gain or loss attributable to each partner must be tracked.

The abstract representation in this thesis attempts to calculate tax liability without regard to the two models. Notation is proposed that is able to produce the correct tax liability given financial activity in which partnerships commonly engage. This allows for an algorithmic approach to calculating tax.

Several non-trivial aspects of partnership activity are not included in the formalism presented below. The most notable are:

- Depreciation deductions
- Partnership liabilities
- Variation in distributive shares of income, gain, loss, liabilities, deductions and credits

Representing such activities using the proposed formalism will be left for future work.

While many secondary sources were used in order to properly interpret the Internal Revenue Code, a majority of the calculations were derived from relevant portions of §§743, 734 and 755 shown in appendix B. Information regarding the calculation of partnership taxation was gathered primarily from the canonical textbook on income taxation of partnerships [11], as well as an online instruction manual intended for Certified Public Accountants [2].

3.1 Formalism

This formalism describes how the IRC dictates that tax accounting should be handled within partnerships. The formalism describes a single partnership, in which an arbitrary number of entities are partners. After its formation, entities can become partners by either *a)* contributing an asset, or *b)* purchasing part or all of a pre-existing partner's share.

Partnership tax accounting can be thought of as a system of entities and assets, where the FMV of the assets is a continuous function, but all other fields are changed only when an *event* occurs. That is, suppose an event occurs at time $t > t' \in \mathbb{R}^+$, and while the event preceding it occurred at time $t' \in \mathbb{R}^+$. Before processing the event, all tax consequences for the partners arising from FMV changes between times t' and t are taken into account. The distance between times is irrelevant, all that matters is that an event must occur.

[2] http://www.cpelite.com/wp-content/uploads/2012/12/PARTNERSHIP-TAXATION-PART-II-2012.pdf

These events include

1. Sale of a partnership interest
2. A distribution, liquidating or non-liquidating
3. Sale of property held by the partnership to a third-party
4. Contribution of an asset by a third-party to gain entry into the partnership

Assumptions

The assumptions listed below are introduced without any loss of generality, and extending the formalism to account for additional cases will be left for future work.

1. An entity has the same share of income as they do liability.
2. All assets are non-depreciable
3. There is only *one* capital gains rate
4. The seven year rules in §§$704(c)(2)(B)$ and $737(b)(1)$ are extended indefinitely.

Assets

An asset is a tuple (b, τ) consisting of

- Adjusted Basis: A scalar $b \in \mathbb{R}^+$
- Type: A positive integer τ that whether the asset is category 0 (cash), category 1 (ordinary) and category 2 (capital).

We can define the set of all assets A as the union of three disjoint subsets $A = A_0 \cup A_1 \cup A_2$, where

- A_0 is the set of all category 0, defined here as any cash or cancellation of debt assets
- A_1 is the set of all category 1 assets, and
- A_2 is the set of all category 2 assets

Entities

An entity is a tuple $(\theta, \kappa, s, \sigma)$ consisting of

- Ordinary Income: A scalar $\theta \in \mathbb{R}$ that records ordinary taxable income for the entity

- Capital Income: A scalar $\kappa \in \mathbb{R}$ that records capital taxable income for the entity
- Share: A scalar $s \in (0, 1)$ that represents the entity's share of partnership income *and* liabilities.[3]
- Outside Basis: A scalar $\sigma \in \mathbb{R}^+$

Partnership Functions

A partnership has three functions that returns information regarding entities and assets. A is the space of all partnership assets and E is the space of all entities that are partners in the partnership

- **FMV** ($F : A \times \mathbb{R}^+ \to \mathbb{R}$): A function $F(a, t)$ that returns the fair market value (FMV) of asset a at time t
- **Built-in Income** ($\Phi : E \times A \to \mathbb{R}$): A function $\Phi(e, a)$ that takes in an entity and an asset and returns the amount of built-in gain or loss allocated to entity e from asset a.
- **Accrued Income** ($\Gamma : E \times A \to \mathbb{R}$): A function $\Gamma(e, a)$ that takes in an entity and an asset and returns the amount of gain or loss allocable to entity e from the change in FMV of asset a

When referring to an element of an entity or asset tuple, we will use subscript notation. For example, to define a variable x as the sum of an asset a's basis and an entity e's share, we would say

$$x = a_b + e_s$$

If we wanted to define another variable y as the difference between the built-in gain or loss of e in respect to a and the *accrued income* of e in respect to a, then we would write

$$y = \Phi(e, a) - \Gamma(e, a)$$

There are times at which the mappings of either the Φ or Γ functions must be changed. For example, if we wanted to indicate that entity e's built-in gain in respect to asset a is now 5, we would write

$$\Phi(e, a) = 5$$

Additionally, there are many instances in which the value of either the Φ or Γ function needs to be incremented by a certain amount. In this case, we introduce a prime notation to represent the new value of the function. For example, if we wanted to indicate that e's accrued income in respect to asset a has increased by 8, then we would write

$$\Gamma'(e, a) = \Gamma(e, a) + 8$$

[3] May later become a vector to represent different allocable shares of income vs liabilities

Technically, this is a contraction of the two lines

$$\Gamma'(e,a) = \Gamma(e,a) + 8$$
$$\Gamma(e,a) = \Gamma'(e,a)$$

but for the remainder of the formalism, we will not include the second line.

Variables are incremented in a similar manner. To show that entity e's capital income has increased by 45, we would write

$$e'_\kappa = e_\kappa + 45$$

3.1.1 Formation of Partnership

When forming a partnership P, an entity can become a partner by either contributing assets or promising to contribute service of some kind. Thus, an entity's initial share in the partnership need not be proportional to FMV of the assets that they contribute.

Thus, we say that a partnership has both a set E and a set A, where $|E| \geq 2$ is a set of entities and $|A| \geq 0$ is a set of assets. The formation of the partnership takes place at $t = 0$. Thus, the adjustment function Φ and accrued income function Γ are initialized to 0, such that $\Phi(e,a) = 0$ and $\Gamma(e,a) = 0$ for all $e \in E$ and $a \in A$.

Each entity is assigned a share $s \in (0,1)$, taking into consideration both the assets that they contributed and any services that they promise to perform. Additionally, both ordinary and capital income is set to zero initially, i.e. $e_\theta = 0$ and $e_\kappa = 0$.

Suppose an entity contributes a set of assets $\mathbf{a}$, as well as possibly offering to contribute future services. Their outside basis is the sum of the adjusted bases of all of the *assets* that they contribute, i.e,

$$e_\sigma = \sum_{a \in \mathbf{a}} a_b \tag{3.1}$$

Additionally, each of these assets may have §704(c) gain or loss, i.e. a difference between its FMV at time of contribution and its adjusted basis. Thus, the built-in function Φ is set such that

$$\Phi(e,a) = \begin{cases} F(a,0) - a_b & \text{if } a \in \mathbf{a} \\ 0 & \text{else} \end{cases} \tag{3.2}$$

3.1.2 Event Update

Recall that immediately before processing an event listed in section 1, tax consequences arising from changes in asset FMVs between the current time t and the time at which the last event occurred $\hat{t}$. Thus, right before each event occurs, we update the accrued income function Γ for all $e \in E$ and $a \in A$ as such

$$\Gamma'(e, a) = \Gamma(e, a) + e_s \cdot \big(F(a, t) - F(a, \hat{t})\big) \tag{3.3}$$

That is, the accrued income of entity e for asset a is incremented by the entity's share multiplied by the amount that the asset has changed in value since time t'

3.1.3 Sale of partnership interest

An entity can sell either all or part of their partnership interest. While there is some confusion in the tax code regarding the tax treatment of partial interest sales when the partnership has liabilities, the following is generally true.

Suppose that entity $\bar{e}$ sells a portion $x \in (0, 1]$ of their partnership interest to a third-party entity e for *cash* at time t. In respect to the sale of a partnership interest, a **substantial built-in loss**, is defined in §743(d):

> For purposes of this section, a partnership has a substantial built-in loss with respect to a transfer of an interest in a partnership if the partnerships adjusted basis in the partnership property exceeds by more than $250,000 the fair market value of such property.

That is, if

$$\sum_{a \in A} (F(a, t) - a_b) \leq -T \tag{3.4}$$

where $T = 250,000$.
Let us now look at the cases where $x = 1$ and $x < 1$

3.1.3.1 Case 1: Entire Share Sold $x = 1$

If $x = 1$, then clearly $\bar{e}$ is selling their entire partnership interest, and following the sale, $\bar{e}$ will be removed from partnership P.

The determination of taxable income begins with determining the difference between the amount realized for the partnership interest and the outside basis of the transferor. Because we assume that FMV of assets takes into account all valuation, we say that the total taxable income I_t arising from the sale is

$$I_t = \bar{e}_s \sum_{a \in A} F(a, t) - \bar{e}_\sigma \tag{3.5}$$

To determine the character of the gain or loss, we refer to §741 of the IRC which states

In the case of a sale or exchange of an interest in a partnership, gain or loss shall be recognized to the transferor partner. Such gain or loss shall be considered as gain or loss from the sale or exchange of a capital asset, except as otherwise provided in section 751 (relating to unrealized receivables and inventory items).

Where §751 defines assets that yield ordinary income in the context of a partnership. Thus, income derived from the increase or decrease in FMV of ordinary assets is taxed as ordinary gain, and the remainder of the total taxable income is capital.

Thus, we must calculate the amount of ordinary income that would be allocated to $\bar{e}$ in the event of a liquidation. That value, I_θ, is calculated as

$$I_\theta = \sum_{a \in A_0 \cup A_1} \Phi(\bar{e}, a) + \Gamma(\bar{e}, a) \tag{3.6}$$

Finally, realized capital income is defined as the difference between total income and ordinary income. Thus, capital income I_κ for $\bar{e}$ is

$$I_\kappa = I_t - I_\theta \tag{3.7}$$

Thus, we increment $\bar{e}_\theta$ by I_θ and $\bar{e}_\kappa$ by I_κ.

The entering partner e takes on an outside basis of whatever s/he paid for the interest and the share of the previous partner, i.e.

$$e_\sigma = \bar{e}_s \sum_{a \in A} F(a, t) \quad \text{and} \quad e_s = \bar{e}_s$$

The next steps depend on whether or not a §743(b) basis adjustment occurs. This occurs if *a)* the partnership had made an optional §754 election or *b)* there is a substantial built-in loss in the partnership's portfolio as defined in equation 3.4.

3.1.3.1.1 No Basis Adjustment Occurs If nothing triggers a §743 basis adjustment, the receiving partner e assumes both the adjustments and accrued income of $\bar{e}$, with a slight stipulation. §704(c)(1)(C) states that

(C) if any property so contributed has a built-in loss

(i) such built-in loss shall be taken into account only in determining the amount of items allocated to the contributing partner, and

(ii) except as provided in regulations, in determining the amount of items allocated to other partners, the basis of the contributed property in the hands of the partnership shall be treated as being equal to its fair market value at the time of contribution.

Thus, because built-in loss can only be claimed by the contributing partner,

$$\Gamma(e,a) = \Gamma(\bar{e},a) \quad \text{and} \quad \Phi(e,a) = \max\{0, \Phi(\bar{e},a)\} \quad \forall a \in A \tag{3.8}$$

3.1.3.1.2 Basis Adjustment Occurs

If either the optional or mandatory triggering of a §743 basis adjustment occurs, then the IRC specifies

> In the case of a transfer of an interest in a partnership by sale or exchange or upon the death of a partner, a partnership with respect to which the election provided in section 754 is in effect or which has a substantial built-in loss immediately after such transfer shall
>
> **(1)** increase the adjusted basis of the partnership property by the excess of the basis to the transferee partner of his interest in the partnership over his proportionate share of the adjusted basis of the partnership property, or
>
> **(2)** decrease the adjusted basis of the partnership property by the excess of the transferee partners proportionate share of the adjusted basis of the partnership property over the basis of his interest in the partnership.

This merely states that unlike the case without basis adjustment, the transferee should *not* inherit any of the transferor's built-in or accrued income.

Thus, rather than assuming the adjustments and accrued income of $\bar{e}$, e maps to 0 for both, so

$$\Phi(e,a) = 0 \quad \text{and} \quad \Gamma(e,a) = 0 \qquad \forall a \in A$$

Next, the adjusted basis of the assets within the partnership must be adjusted further to reflect previously realized gains or losses by $\bar{e}$. Thus, all adjusted bases in the partnership are incremented as such.

$$a'_b = a_b - \Phi(\bar{e},a) - \Gamma(\bar{e},a) \quad \forall a \in A \tag{3.9}$$

3.1.3.2 Case 2: Part of Share Sold

$x < 1$

If $x < 1$, then $\bar{e}$ is selling only part of their share. Thus, the calculation of realized income is nearly identical to those shown in equations 3.5, 3.6 and 3.7, except multiplied by the fraction of the share being sold. That is,

$$\bar{e}'_\sigma = \bar{e}_\sigma + I_\theta \quad \text{where } I_\theta = x\Big(\sum_{a \in A_0 \cup A_1} \Phi(\bar{e},a) + \Gamma(\bar{e},a)\Big) \tag{3.10}$$

$$\bar{e}'_\kappa = \bar{e}_\kappa + I_t - I_\theta \quad \text{where } I_t = x\Big(\bar{e}_s \cdot \sum_{a \in A} F(a,t) - \bar{e}_\sigma\Big) \tag{3.11}$$

Similarly, e takes on a fraction of both the outside basis and share that they would have received if they purchased the entire share, i.e.

$$e_\sigma = x\bar{e}_s \sum_{a \in A} F(a,t) \quad \text{and} \quad e_s = x * \bar{e}_s$$

3.1.3.2.1 No Basis Adjustment Occurs If there is no §743 basis adjustment, e's adjustment and accrued income functions follow a similar pattern as in equation 3.8

$$\Gamma(e,a) = x\Gamma(\bar{e},a) \tag{3.12}$$

$$\Phi(e,a) = x\max\{0, \Phi(\bar{e},a)\} \quad \forall a \in A \tag{3.13}$$

Unlike the whole share sale, we now need to update $\bar{e}$'s adjustment and accrued income function for all $a \in A$ as such

$$\Gamma(\bar{e},a) = (1-x)\Gamma(\bar{e},a) \tag{3.14}$$

$$\Phi(\bar{e},a) = (1-x)\max\{0, \Phi(\bar{e},a)\} \tag{3.15}$$

This accounts for the fact that there still remains a $(1-x)$ portion of $\bar{e}$'s allocable income that hasn't been allocated yet.

3.1.3.2.2 Basis Adjustment Occurs If a basis adjustment as described in the previous section occurs, then as with the whole interest sale, both the adjustment and accrued interest function of e map to zero for every asset. Now increment the adjusted bases of all of the assets in the partnership as such

$$a'_b = a_b - x\Big(\Phi(\bar{e},a) + \Gamma(\bar{e},a)\Big) \quad \forall a \in A \tag{3.16}$$

After this is done, complete the same adjustment to the Φ and Γ functions described in equation 3.14.

3.1.4 Distributions

Distributions are technically a transaction between the partnership and a partner, in which the partnership gives the partner assets in exchange for all of part of their share. The shares of all of the other partners are subsequently changed, due to the change in the FMV of the partnership.

Definition Given an asset a and a set of assets $\hat{A}$ at time t

The *total value increase* function

$$\Delta F_+(\hat{A}) = \sum_{\breve{a} \in \hat{A}} \max\{F(\breve{a}, t) - \breve{a}_b, 0\} \tag{3.17}$$

The *proportional value increase* function

$$W_+(a, \hat{A}) = \frac{\max\{F(a, t) - a_b, 0\}}{\Delta F_+(\hat{A})} \tag{3.18}$$

The *total value decline* function

$$\Delta F_-(\hat{A}) = \sum_{\breve{a} \in \hat{A}} \max\{\breve{a}_b - F(\breve{a}, t), 0\} \tag{3.19}$$

The *proportional value decline* function

$$W_-(a, \hat{A}) = \frac{\max\{a_b - F(a, t), 0\}}{\Delta F_-(\hat{A})} \tag{3.20}$$

The *proportional FMV* function

$$\begin{aligned} W(\hat{A}) &= \frac{F(a, t)}{\bar{a}}, \text{ where} \\ \bar{a} &= \sum_{\breve{a} \in \hat{A}} F(\breve{a}, t)) \end{aligned} \tag{3.21}$$

The *proportional basis* function

$$\begin{aligned} W_B(\hat{A}) &= \frac{a_b}{\tilde{a}}, \text{ where} \\ \tilde{a} &= \sum_{\breve{a} \in \hat{A}} \breve{a}_b \end{aligned} \tag{3.22}$$

Suppose at time t, partner e receives a set of assets $\mathbf{a}$ as a distribution.

Definition A *liquidating distribution* occurs when the sum of the FMVs of the distributed assets are within an acceptable range of the distributee's share of the total partnership assset FMVs, i.e.

$$\sum_{a \in \mathbf{a}} F(a, t) = e_s \cdot \sum_{a \in A} F(a, t) \pm \epsilon \tag{3.23}$$

where $\epsilon \in \mathbb{R}^+$ represents the fact that distribution of assets may not be *exactly* the value of the entity's share.

Definition Conversely, a *non-liquidating distribution* occurs when

$$\sum_{a\in\mathbf{a}} F(a,t) < e_s \cdot \sum_{a\in A} F(a,t) - \epsilon \tag{3.24}$$

because the combined FMV of the received assets is significantly less than the value of e's current interest.

Subject to §704(c)(1)(B) if any of the distributed assets were originally contributed to the partnership by any of the existing partners, then any gain or loss attributable to them must be realized, which includes incrementing their outside basis. So for each distributed asset $a \in \mathbf{a}$, and each $e \in E \setminus e$, increment the appropriate income class.

$$e'_\theta = e_\theta + \Phi(e,a) \quad \text{if } a \in A_1 \tag{3.25}$$

$$e'_\kappa = e_\kappa + \Phi(e,a) \quad \text{if } a \in A_2 \tag{3.26}$$

and adjust the entity's outside basis, as well as the adjusted basis for the asset.

$$e'_\sigma = e_\sigma + \phi(e,a) \tag{3.27}$$

$$a'_b = a_b + \Phi(e,a) \tag{3.28}$$

Definition For clarification during the remainder of the section, the following variables will be used.

$$\mathbf{a}_0 = \mathbf{a} \cap A_0 \quad \alpha_0 = \sum_{a\in\mathbf{a}_0} F(a,t) \tag{3.29}$$

$$\mathbf{a}_1 = \mathbf{a} \cap A_1 \quad \alpha_1 = \sum_{a\in\mathbf{a}_1} a_b \tag{3.30}$$

$$\mathbf{a}_2 = \mathbf{a} \cap A_2 \quad \alpha_2 = \sum_{a\in\mathbf{a}_2} a_b \tag{3.31}$$

We now look at cases regarding distributions, which are differentiated primarily by the distributee's outside basis, as well as the character of the bases of the distributed assets.

3.1.4.1 Outside Basis Exceeds Sum of Distributed Bases

$$e_\sigma > (\alpha_0 + \alpha_1 + \alpha_2)$$

If the distribution is non-liquidating, then no taxable income is allocated to e. The distributee receives all assets, and the basis of each asset a^i remains the most recent a^i_b value that it held. e's new outside basis becomes its previous outside basis minus the combined bases of the assets it received, i.e.

$$e'_\sigma = e_\sigma - \sum_{a \in \mathbf{a}} a_b \tag{3.32}$$

If this is a liquidating distribution, the problem is that e has some outside basis left in the partnership that s/he must allocate amongst the bases of the assets. The bases of the category 1 assets cannot be stepped-up above what they already are, so all of the excess basis must be allocated amongst the category 2 assets. This is a two-step process, which begins with trying to bring the basis of both capital assets up to their FMV.

Definition The excess basis

$$B = e_\sigma - \sum_{a \in \mathbf{a}} a_b \tag{3.33}$$

3.1.4.1.1 Excess Basis Less than Category 2 Appreciation

$$B \leq \Delta F_+(\mathbf{a}_2)$$

That is, if the sum of the positive differences between all category 2 FMVs and their respective adjusted basis, is less than (or equal to) the excess basis. In this case, there is not (or just) enough outside basis to increase the adjusted basis of the assets to their FMV.

Here, the excess basis B is allocated to all category 2 assets proportional to the positive difference between their FMV and adjusted basis. That is, for $a \in \mathbf{a}_2$

$$a'_b = a_b + B \cdot W_+(a, \mathbf{a}_2) \tag{3.34}$$

3.1.4.1.2 Excess Basis Greater than Category 2 Appreciation

$$B > \Delta F_+(\mathbf{a}_2)$$

That is, if there is more than enough excess outside basis to increase the adjusted bases of all of the distributed category 2 assets up to their FMVs, then something must be done with the additional excess outside basis.

Definition The amount of outside basis that still needs to be allocated to category 2 asset adjusted bases is

$$\hat{B} = B - \Delta F_+(\mathbf{a}_2) \tag{3.35}$$

Similar to the case described in section 3.1.4.1.1, the additional outside basis is allocated to each adjusted basis proportional to their FMV. That is, the adjusted basis of each $a \in \mathbf{a}_2$ becomes

$$a'_b = \max\{a_b, F(a,t)\} + \hat{B} \cdot W(a, \mathbf{a}_2) \tag{3.36}$$

The first term shows that if the adjusted basis was previously less than the asset's FMV, then we add to the FMV. Conversely, we add to the adjusted basis if it was higher than the FMV to begin with. The second term adds the amount of the additional excess basis proportional to the FMV of each asset

3.1.4.2 Outside Basis Exceeds Cash and Category 1 Basis

$$e_\sigma > \alpha_0 + \alpha_1$$

Here we cover the situation where the amount of outside basis exceeds the necessary amount to cover both category 0 and category 1 assets, but not any category 2 assets that may be distributed. If no category 2 assets are distributed *and* the distribution is *liquidating*, the distributee claims an ordinary loss and nothing else occurs, i.e.

$$e'_\theta = e_\theta - \big(e_\sigma - (\alpha_0 + \alpha_1)\big) \quad \text{if } \mathbf{a}_2 = \emptyset \tag{3.37}$$

but that does not leave enough outside basis to compensate for the adjusted bases of all of the category 2 assets. So the allocation is similar to that in section 3.1.4.1

The difference is that instead of adding to the adjusted bases of the category 2 assets, we subtract from them.

Definition The amount that we must subtract amongst all category 2 asset bases is

$$D_2 = \Big(\sum_{a \in \mathbf{a}_2} a_b\Big) - \big(e_\sigma - (\alpha_0 + \alpha_1)\big) \tag{3.38}$$

3.1.4.2.1 Basis Shortfall less than Category 2 Value Decline

$$D_2 \leq \Delta F_-(\mathbf{a}_2)$$

In this case, the distributed category 2 assets have declined in value since they have been in the partnership *more* than the amount that must be subtracted from all of their adjusted bases. Thus, we just subtract the bases proportional to the amount, if any, that each distributed category 2 asset has declined in value. Thus, each adjusted basis for each $a \in \mathbf{a}_2$ becomes

$$a'_b = a_b - D_2 \cdot W_-(a, \mathbf{a}_2) \tag{3.39}$$

In other words, we subtract from the adjusted basis of each asset proportional to the *positive* difference between their adjusted basis and FMV.

3.1.4.2.2 Basis Shortfall greater than Category 2 Value Decline

$$D_2 > \Delta F_-(\mathbf{a}_2)$$

If this is the case, then even after accounting for decline in category 2 asset value, more basis must be subtracted from the assets' adjusted bases.

Definition The additional basis decrease is calculated as

$$\hat{D}_2 = D_2 - \Delta F_-(\mathbf{a}_2) \tag{3.40}$$

to be the amount of basis that still needs to be subtracted from the distributed category 2 bases in a manner subject to §755, which specifies that the residual outside basis subtraction take place proportional to the assets' FMVs. The adjusted bases for every asset $a \in \mathbf{a}_2$ thus becomes

$$a'_b = \min\{a_b, F(a,t)\} - \hat{D}_2 \cdot W(a, \mathbf{a}_2) \tag{3.41}$$

3.1.4.3 Outside Basis Exceeds Cash Value

$$e_\sigma < \alpha_0 + \alpha_1$$

In this case, §755 dictates that the adjusted basis of the capital assets in the hands of the distributee is zero, i.e.

$$a_b = 0 \quad \forall a \in \mathbf{a}_2 \tag{3.42}$$

Implicitly, this means that the D_2 variable defined in equation 3.38 is the sum of the distributed category 2 assets, i.e.

$$D_2 = \sum_{\mathbf{a}_2} a_b \tag{3.43}$$

The next step is to determine how to allocate the outside basis to the category 1 assets, and if the transaction is taxable.

3.1.4.3.1 Case 1: $e_\sigma > \alpha_0$

In this case, e has greater outside basis than cash or cancellation of debt, but there is not enough outside basis to allocate fully to the adjusted bases of all distributed category 1 assets. Thus, using the same heuristic to detract from adjusted basis as shown in equations 3.39 and 3.41, we perform the same steps on category 1 assets.

Definition The amount that needs to be subtracted collectively from all distributed category 1 assets

$$D_1 = \Big(\sum_{a \in \mathbf{a}_1} a_b\Big) - (e_\sigma - \alpha_0) \tag{3.44}$$

if $D_1 \leq \Delta F_-(\mathbf{a}_1)$, then we subtract from the adjusted bases of the distributed category 1 assets proportional to their depreciation as in equation 3.39. That is, the adjusted basis of every asset $a \in \mathbf{a}_1$ becomes

$$a'_b = a_b - D_1 \cdot W_-(a, \mathbf{a}_1) \tag{3.45}$$

If $D_1 > \Delta F_-(\mathbf{a}_1)$, then we subtract from adjusted basis in a manner similar to 3.41. Defining the additional excess basis that must be subtracted as $\hat{D}_1 = D_1 - \Delta F_-(\mathbf{a}_1)$, the adjusted basis of every asset $a \in \mathbf{a}_1$ becomes

$$a_b = \min\{a_b, F(a,t)\} - \hat{D}_1 \cdot W(a, \mathbf{a}_1) \tag{3.46}$$

3.1.4.3.2 Case 2: $e_\sigma \leq \alpha_0$ The final case occurs when the distributee entity e has an outside basis that is less than or equal to the amount of cash and cancellation of debt (A_0 assets) that is distributed to it. Here, e receives ordinary income of the difference between the sum of the FMVs of all distributed A_0 assets and e's outside basis. Additionally, all of the distributed category 1 and 2 assets have an adjusted basis of 0 in the hands of e. i.e.

$$a_b = 0 \quad \forall a \in A_1 \cap A_2 \tag{3.47}$$

$$e'_\theta = e_\theta + (e_\sigma - \alpha_0) \tag{3.48}$$

Because the basis of all distributed category 1 and category 2 assets must be adjusted down to zero, we set the basis adjustment variables to reflect that

$$D_2 = \Delta F_-(\mathbf{a}_2) \tag{3.49}$$

$$D_1 = \Delta F_-(\mathbf{a}_1) \tag{3.50}$$

3.1.4.4 §734 Basis Adjustments

Similar to §743 adjustments as described in section 3.1.3, a §734 basis adjustment is necessary if *a)* the partnership has made a §754 election or *b)* there is a substantial built-in loss

However, the definition of "substantial built-in loss" in respect to §734 is different than in respect to §743.

Definition In respect to §734, a "substantial built-in loss" is when

$$-(\Lambda_1 + \Lambda_2) \geq T$$

where like the §743 adjustment, $T = \$250,000$ as seen in §734$(d)(1)$, and Λ_1 and Λ_2 refer to, respectively, category 1 and category 2 basis adjustments

There are two events that trigger a §734 basis adjustment.

1. The first event is the realization of any gain or loss by the distributee upon distribution, the realizations are mutually exclusive.

   - **Gain** is realized to the extent that cash or cancellation of debt exceeds the distributee's outside basis as in section 3.1.4.3.2, i.e.

   $$\Lambda_2 = \max\{\alpha_0 - e_\sigma, 0\}$$

   - **Loss** is realized *only* in liquidating distributions, as shown in section 3.1.4.2. Thus, the amount of loss realized is

   $$\begin{aligned} \Lambda_2 = \max\{e_\sigma - (\alpha_0 + \alpha_1), 0\} \\ \text{if} \quad \mathbf{a}_2 = \emptyset \end{aligned} \tag{3.51}$$

2. The second of event that triggers a §734 basis adjustment is the upward or downward adjustment of a distributed assets' basis. As we've seen in this section, this occurs in all cases other than (add reference later). Upon any adjustment of the distributed assets' bases, an adjustment occurs on the bases of *similar* typed assets in the *opposite* direction from within the partnership.

   Recall B, D_2 and D_1 as described in, respectively, equations 3.33, 3.38 and 3.44, the basis adjustments are further incremented as such

   $$\Lambda_1' = \Lambda_1 + D_1 \tag{3.52}$$

   $$\Lambda_2' = \Lambda_2 + D_2 - B \tag{3.53}$$

If either Λ_1 or Λ_2 is zero, then there is no basis adjustment within the respective category of assets.

Upon determination of the Λ_1 and Λ_2 variables, they are applied to the adjusted basis of the assets within the partnership as described below. Post-distribution, all of the distributed assets $a \in \mathbf{a}$ have been removed from the partnership and thus are not affected by the following basis adjustments. The process requires adjusting the bases of each category separately. The remainder of this section is applied first for $Q = 1$, then $Q = 2$ to illustrate the basis adjustments taking place within each asset category.

Definition Either the total increase or decrease of FMV amongst assets within the partnership is defined as

$$\Delta F(A_Q, t) = \begin{cases} \Delta F_+(A_Q) & \text{if } \Lambda_Q > 0 \\ \Delta F_-(A_Q) & \text{if } \Lambda_Q < 0 \end{cases} \tag{3.54}$$

3.1.4.4.1 Case 1: Asset FMV Change Less than Adjustment

$\Delta F(a,t) \leq |\Lambda_Q|$ Here we use the adjustment to fully or partially close the gap between certain assets' FMV and adjusted basis. If $\Lambda_Q > 0$, then the gap is closed for assets that have *appreciated* in value because the positive adjustment can bring the adjusted basis closer to its FMV. By the same logic, we close the gap for assets that have *declined* in value when $\Lambda_Q < 0$, i.e., for each asset $a \in A_Q$

$$a_b' = \begin{cases} a_b + \Lambda_Q \cdot W_+(a, A_Q) & \text{if } \Lambda_Q > 0 \\ a_b + \Lambda_Q \cdot W_-(a, A_Q) & \text{if } \Lambda_Q < 0 \end{cases} \tag{3.55}$$

The *accrued income* function for each entity must be adjusted in an equivalent manner to reflect the basis adjustment. That is for each $e \in E$ and $a \in A_Q$,

$$\Gamma'(e,a) = \begin{cases} \Gamma(e,a) - e_s \cdot A_Q \cdot W_+(a, A_Q) & \text{if } \Lambda_Q > 0 \\ \Gamma(e,a) - e_s \cdot A_Q \cdot W_-(a, A_Q) & \text{if } \Lambda_Q < 0 \end{cases} \tag{3.56}$$

3.1.4.4.2 Case 2: Asset FMV Change Greater than Adjustment

$\Delta F(A_Q t) > |\Lambda_Q|$ Here, there is enough positive (or negative) adjustment to close the positive (or negative) gap between each asset's FMV (or basis) and adjusted basis. Thus, the additional adjustment must be applied proportional to the FMVs of all assets in the appropriate category.

Definition The additional basis adjustment is

$$\hat{\Lambda}_Q = \begin{cases} \Lambda_Q - \sum_{a \in A_Q} \max\{F(a,t) - a_b, 0\} & \text{if } \Lambda_Q > 0 \\ \Lambda_Q + \sum_{a \in A_Q} \max\{a_b - F(a,t), 0\} & \text{if } \Lambda_Q < 0 \end{cases} \tag{3.57}$$

If $\hat{\Lambda}_Q > 0$, then each the adjusted basis of each asset $a \in A_Q$ is increased as such

$$a_b' = \max\{a_b, F(a,t)\} + \hat{\Lambda}_Q \cdot W(a, A_Q) \tag{3.58}$$

Alternatively, if $\hat{\Lambda}_Q < 0$, then

Definition The *negative allocation* that each asset a is allocated is

$$\mu = \min\{a_b, F(a,t)\} + \hat{\Lambda}_Q \cdot W_B(a, A_Q) \tag{3.59}$$

Thus, the adjusted bases of each asset $a \in A_Q$ are decreased as such

$$a'_b = a_b = \max\{\mu, 0\} \tag{3.60}$$

Definition There is an additional

$$\breve{\Lambda}_Q = \sum_{a \in A_Q} \min\{\mu, 0\} \tag{3.61}$$

negative basis adjustment that is suspended until another asset from the associated category Q enters the partnership, at which point the same process from equations 3.57 and 3.60 until $\breve{\Lambda}_Q$ is exhausted. If $\breve{\Lambda}_Q = 0$, then the process is finished.

Again, we must change each partner's *accrued income* function to reflect the basis adjustment. For each $e \in E$ and $a \in A_Q$

$$\Gamma'(e, a) = \begin{cases} \Gamma(e, a) - e_s\Big(\max\{F(a, t) - a_b, 0\} + \hat{\Lambda}_Q \cdot W(a, A_Q)\Big) & \text{if } \Lambda_Q > 0 \\ \Gamma(e, a) + e_s\Big(\min\{F(a, t) - a_b, 0\} + \hat{\Lambda}_Q \cdot W_B(a, A_Q)\Big) & \text{if } \Lambda_Q < 0 \end{cases} \tag{3.62}$$

3.1.5 Sale of Asset

Suppose at time t, the partnership sells some category 1 or 2 asset $a \in A$ for its FMV. For each entity $e \in E$, increment their taxable incomes as such

$$\begin{aligned} e'_\theta &= e_\theta + \Phi(e, a) + \Gamma(a, e) \quad \text{if } a \in A_1 \\ e'_\kappa &= e_\kappa + \Phi(e, a) + \Gamma(a, e) \quad \text{if } a \in A_2 \end{aligned} \tag{3.63}$$

4

Simulation

The following section answers the second question proposed in section 1.2.1.

> How can we simulate the process of an IRS auditor, under resource constraints, determining the aspects of a sequence of transactions which indicate suspicious behavior?

Crucial to the crafting of better auditing procedures is a means of simulating *a)* the initial network of taxpayers and partnerships that is being evaluated, *b)* a sequence of transactions occurring within that network, and *c)* a representation of an auditing procedure that assigns the likelihood of an audit to a transaction sequence and initial network. The former two representations are loosely based on the formalism described in section 3.1. Due to the computational nature of the method described in this thesis, a numerical representation of the way in which auditors assign risk to financial activity is required.

4.1 Tax Network

As previously discussed in chapter 3, this thesis focuses on partnership structures, which are generally a set of interconnected taxpayers and partnerships. See figure 4.1 below for an example of a tax network.

Here, ownership of an interest in a partnership is represented by a straight arrow from the owner to the partnership. The number in parenthesis displays the owner's outside basis in its share of the partnership. In figure 4.1, Alice owns a 50% share in $P1$ and an outside basis of \$40. Similarly, Bob owns 30% with an outside basis of \$30, and Cathy owns a 20% share with an outside basis of \$18. Additionally, the entities all have a portfolio of assets, each of which has an FMV and adjusted basis [1]. The taxpayers' portfolios are not shown in figure 4.1 for clarity, but they have the same structure as the partnership's portfolio.

[1] The adjusted basis of cash is always equal to its FMV, so it is not included in the portfolio

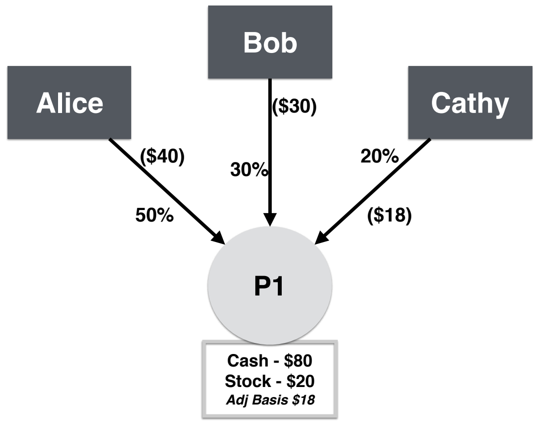

Figure 4.1: Example of a network of interconnected entities. The number in parenthesis denotes each owner's outside basis.

In this way, the simulation is primed with a set of initial conditions, which are specified by the portfolios of the entities. That is, each entity's interest in a partnership is an asset in their portfolio, thus the initial economic position of the network is described by the assets held by the entities.

4.2 Transaction Sequence

Transactions are implemented on a tax network in a manner that exchanges assets between entities. As seen below in figure 4.2, Emma purchases Alice's 50% interest in $P1$ for \$50, as represented by the dotted line. Upon purchase, Emma's outside basis in $P1$ becomes \$50.

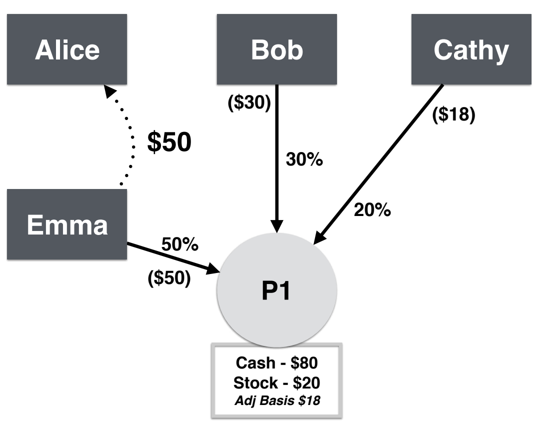

Figure 4.2: Example of a transaction between entities. The dotted line denotes the transfer of assets.

Each transaction, if legal, changes the state of the network by transferring assets between entities. In figure 4.2, \$50 that used to be in Emma's portfolio is now in Alice's, while Emma owns the interest in $P1$ that used to be owned by Alice. A transaction sequence can thus be viewed as a set of operators that transitions the network from one state to another.

4.2.1 Final Tax Network State

Restrictions can be placed on transaction sequences such that the final network state meets certain requirements. For example, suppose that in the scenario described by figures 4.1 and 4.2, the requirement was that in the final network state, Emma must own 50% of $P1$. In that case, Alice could have instead sold her share to Emma in exchange for an annuity or stock worth $50. Alternatively, Emma could have engaged in two separate transactions, purchasing Bob's 30% share for $30 and subsequently purchasing Cathy's 20% share for $20. All of these alternate transaction sequences result in Emma owning 50% of $P1$, as seen in figure 4.2.

In this way, two transaction sequences can produce in the same (or similar) economic result, but produce wildly different tax consequences. When Emma purchases Alice's share in cash, Alice is allocated $10 in taxable income because her outside basis in the partnership interest is $40. If Emma purchased the share with an annuity, then the sale would result in no current taxable income for Alice [2]. This conception of restricting transaction sequences to produce a certain tax network state will come more into play in chapter 5.

4.3 Audit Score Sheets

The purpose of an *audit score sheet* is to determine the likelihood of conducting an audit of a partnership tax return, which is a set of information regarding a set of taxpayers and entities, along with a sequence of transactions between them. Thus, the *audit score* is calculated as the sum of all of the *audit points*, each multiplied by the frequency that its associated observable event occurs in the transaction sequence. That *audit score* is the metric used to represent the relative likelihood that a certain tax return should be audited by the IRS.

An IRS auditor is represented as a list of "observable" events, each with an associated positive real number between 0 and 1, which are called *audit points*. In order to simulate resource constraints, it is required that the audit points sum to exactly 1. The list of all observable events with their associated *audit points* is referred to as an *audit score sheet*. One can imagine a hypothetical auditor scanning a set of financial documents and noting when an observable event is present.

The events on the audit score sheet can range from basic facts about a transaction, such as whether a stock is being exchanged, to more complex aspects of the partnership structure state, such as ownership linkages between multiple entities. An observable event

[2] Alice would eventually be allocated taxable income depending on the stipulations of the annuity. The specifics of annuity taxation is beyond the scope of this thesis.

is one that is possible to detect in the tax ecosystem model, but not necessarily by the IRS. For example, if a taxpayer purchases a share in a partnership for cash, the simulation will process that as a transaction involving a partnership asset, as well as tracking all parties involved in the transaction, while in reality, that information might not be accessible in the financial documentation. Three possible types of events are observable within simulation. *1)* Events that are possible to observe *2)* Events that may be impossible or very difficult to observe, but can be captured with an observable proxy *3)* Events that are impossible to observe due to either logistical or legal reasons.

Inspiration for this modeling approach was gathered from the way in which the IRS in the past has amended the tax code to prevent what are believed to be abusive tax shelters. For example, in 2004 the IRC was altered in §743(a) to read

> The basis of partnership property shall not be adjusted as the result of (1) a transfer of an interest in a partnership by sale or exchange or on the death of a partner unless (2) the election provided by §754 (relating to optional adjustment to basis of partnership property) is in effect with respect to such partnership or (3) unless the partnership has a substantial built-in loss immediately after such transfer.

We have added numbers in parenthesis to signify *observable* events: (1) The sale of a partnership interest in exchange for a *taxable* asset. (2) The partnership whose shares are being transferred has not made a §754 election. (3) The seller's basis in respect to the non-cash assets owned by the partnership exceeds their FMV by more than $250,000.

Supposing that the three events above were the only observable events, the template for an *audit score sheet* would be as shown in table 4.1 below.

Observable	**Points**	**Frequency**
1	$Point_1$	$Frequency_1$
2	$Point_2$	$Frequency_2$
3	$Point_3$	$Frequency_3$
$1 \cup 2$	$Point_{1\cup2}$	$Frequency_{1\cup2}$
$1 \cup 3$	$Point_{1\cup3}$	$Frequency_{1\cup3}$
$2 \cup 3$	$Point_{2\cup3}$	$Frequency_{2\cup3}$
$1 \cup 2 \cup 3$	$Point_{1\cup2\cup3}$	$Frequency_{1\cup2\cup3}$

Table 4.1: Each row has three columns with 1) the type of observable corresponding to the three characterized observables from the IRS notice, 2) the associated audit point and 3) the number of times it occurs in a list of transactions

Not only is the occurrence of any of the individual events a row on the *audit score sheet*, but all possible combination of their occurrences are assigned an *audit point* as well. This allows for an auditor to ignore the occurrence of common events, such as the simple sale of a partnership interest, while being able to encapsulate scenarios in which they may indicate abuse or non-compliance. Thus, if there are m individual observable events, then there will be $2^m - 1$ elements on an *audit score sheet*, each of which is assigned an *audit point*. This example will be discussed further in chapter 6.

Thus, given a *a)* tax network of taxpayers and other entities, *b)* sequence of transaction between them [3], and *c)* an *audit score sheet* with *audit points*, we can calculate both the taxable income generate from the financial activity and an *audit score* that represents the likelihood of an audit as shown in figure 4.3.

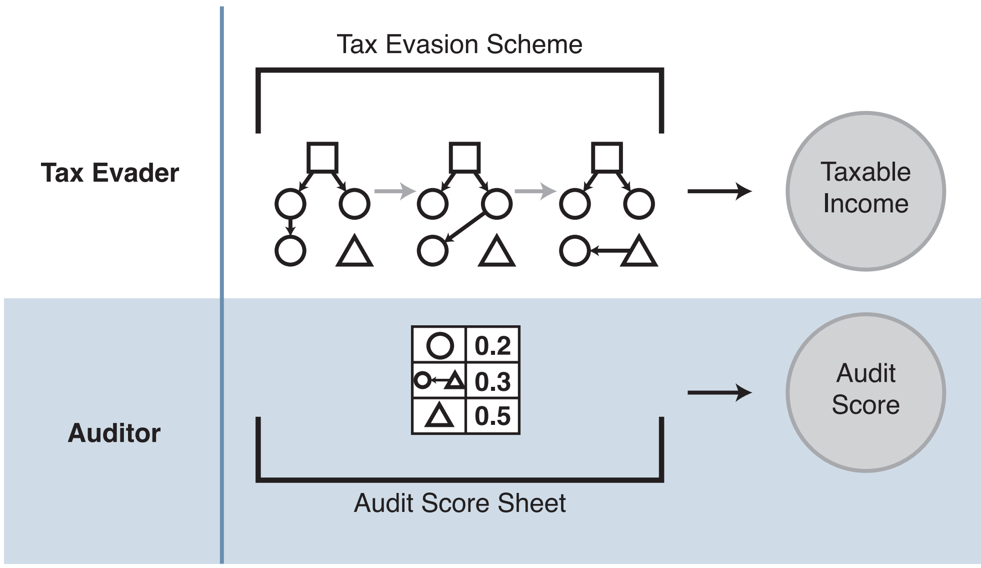

Figure 4.3: A transaction sequence alters the state of the tax network, as well as producing both taxable income and the likelihood of an audit

4.4 Description of Transactions and Auditing

Formalizing the simulation process requires an abstract representation of *a)* a *tax network* of taxpayers and partnerships, all of which have asset portfolios, *b)* a *transaction sequence* that is imposed on the tax network and exchanges assets between members of the network, and *c)* an *audit score sheet* that contains information regarding what types of activities that an auditor finds suspicious. The simulation described below has been implemented in Java as a proof of concept.

[3]with the calculations specified in section 3.1

4.4.1 Tax Network

The tax network at a given time is defined as a list of entities (a term which we define to encapsulate taxpayers and partnerships), each of which owns a set of assets. At any point, the state of the network can be described as some $\gamma \in \Gamma$, where $\gamma = \{\mathbf{e}, \mathbf{a}, d\}$, where $\mathbf{e} = \{e_i\}_{i=0}^{k_1}$ is the set of entities, $\mathbf{a} = \{a_i\}_{i=0}^{k_2}$ is the set of all assets and $k_1, k_2 \in \mathbb{Z}_+, e_i \in E, a_i \in A$. The operator d determines the owner of each asset, i.e $d : A \mapsto E$, where A is the space of assets and E is the space of entities.

4.4.2 Transaction Sequence

We can define a sequence of transactions as a vector $\mathbf{t} = \{t_i\}_{i=0}^{k}$ for some $k \in \mathbb{Z}_+$, $t \in \mathbf{T}$ is the space of all transactions. A transaction is defined as $t = \{e_f, e_t, a_f, a_t\}$, where $e_f, e_t \in E$ are two entities and $a_f, a_t \in A$ are two assets that are being exchanged between the two entities.

4.4.3 Audit Score Sheet

Suppose that there are m specific types of events that are observable, represented by $\{b_i\}_{i=0}^{n}$, where $n = 2^m - 1$. Associated with each type of event are the audit points $\{\alpha_i\}_{i=0}^{n}, \alpha \in \mathbb{R}$ and the frequency that the event occurs within a network of transactions $\{f_i\}_{i=0}^{n}, f_i \in \mathbb{Z}_+$. We can then write the audit score, s corresponding to the audit score sheet and network of transactions as

$$s = \sum_{i=0}^{n} \alpha_i f_i \text{ where } \sum_{i=0}^{n} \alpha_i = 1$$

Legality Representation

We observe that laws governing a given transaction depend on the "type" of assets and entities being exchanged. For example, the laws governing the exchange of a hotel for cash between two taxpayers are different from those governing the contribution of an annuity to a partnership in exchange for a share. Thus, we can determine the laws governing a given transaction by the combination of both asset and entity types.

Consider the transaction $t = (e_f, e_t, a_f, a_t)$, which states that entity e_f gives e_t the asset a_f in exchange for a_t. Define $\hat{E}$ to be the finite set of entity *types*, and $\hat{A}$ to be the finite set of asset *types*. We can then write the set of all transactions as a union of disjoint subsets $\mathbf{T} = \cup_{i=0}^{n} T_i$, where each subset contains all transactions of a certain combination of asset and entity types. The steps that follow are.

1. a transaction type t is first checked to see if it is within the bounds of the legal/feasible region by first determining to which subset T_i it belongs. We define $\mu : T_i \mapsto \Phi$ as a

map from a subset $T_i \in \mathbf{T}$ to Φ that determines the laws ϕ that govern the transaction, given its combination of asset/entity types.

2. the transfers in the two actions composing the transaction represent the transition of the network state γ_t to γ_{t+1} and γ_{t+1} to γ_{t+2} according to the map $\tau : \mathbf{T} \times \Gamma \mapsto \Gamma$

3. taxable income calculation takes a transaction t and a network state γ_t and maps it to a taxable income value η for each taxable entity and an updated network state, $P : \mathbf{T} \times \Gamma \mapsto \mathbb{R} \times \Gamma$

4.4.4 Simulation Description

A transaction sequence is applied to the tax network to produce total *taxable income* of all entities, and the audit score sheet produces an *audit score* in respect to the transaction sequence and tax network.

Thus, the simulation is defined as a function $\mathbf{F} : \mathbf{T} \times \Gamma \times \Psi \mapsto \mathbb{R}_+^2$ that takes as input a sequence of transactions, an initial tax network state and audit score sheet, and generates taxable income and an audit score.

5

Optimization

The final section of the method presented in this thesis confronts the questions posed in section 1.2.1, as well as by multiple tax academics [35, 26]:

> Can new forms of abusive tax behavior be anticipated by optimizing both tax planning and auditing policy?
>
> - How can transaction sequences be constructed which minimize tax liability, while attracting minimal IRS suspicion?
> - How can IRS policy be formed which assigns high auditing likelihood to suspicious transaction sequences, without falsely auditing routine transactions.

Using the language developed in chapters 3 and 4, we can rephrase this as two analogous questions.

5.1 Transaction Sequence Optimization

Given an initial tax network, an *audit score sheet* and requirements regarding the final tax network state, what sequence of transactions will minimize both taxable income and the *audit score*?

As discussed in section 4.2.1, different transaction sequences that produce the same economic result can generate very different tax consequences. Prior work by Howard Abrams shows that different methods of exiting a partnership can be strategic for both the exiting and existing partners [2]. In a simplified example, a partner could sell their interest in a partnership or take a liquidating distribution and incur a certain tax liability. Alternatively, they could (1) take out a large loan through the partnership to increase their share of liabilities, then (2) sell a majority of their interest to a third party, while maintaining a negligible share in the partnership. This way, the partner receives close to the entire value of their interest in cash, while incurring a significantly smaller amount of tax.

Thus, assuming that there is a near infinite space of transaction sequences that produce the same (or similar) economic results, it is the goal of the tax planner to choose the "best" sequence by searching over all combinations of transaction seqyences. This method assumes that a transaction sequence is of high value if it produces low levels of tax liability without arising suspicion from the IRS, i.e. generating a high *audit score*.

5.2 *Audit Score Sheet* Optimization

Given an initial tax network, a transaction sequence and a required final tax network state, what vector of *audit points* will generate a high *audit score* when taxable income is low relative to other transaction sequences that produce the same economic result?

The optimization of an *audit score sheet* is essentially a method for classifying an abusive tax shelter. That is, the goal is to find the joint occurrence of *observable* events which perfectly describe the shelter, such as the three example *observables* discussed in section 4.3. Not only will this result in every similarly abusive transaction sequence to be audited, but will prevent against "false positive" audits on legitimate transaction sequences.

Specifically, we assume that the value of an *audit score sheet* is based on a combination of the generated tax liability and the *audit score*. If a transaction sequence results in a low tax liability relative to other transaction sequences that produce similar economic results, then the *audit score* should be high. Conversely, average or high levels of tax liability should be associated with a low *audit score* so as not to waste auditing resources on routine, non-abusive transactions.

5.3 Evolving Transactions and Audits

The goal laid out in sections 5.1 and 5.2 requires a method to search over a near infinite space of potential solutions, otherwise known as a *search heuristic*. This thesis uses a class of *search heuristics* known as Evolutionary Algorithms (EA). Due to previous work suggesting that tax planners and IRS policies co-evolve with one another [35, 26], an algorithm based on neo-Darwinian concepts is a natural method for finding optimal solutions [33, 25, 16, 4] [1].

EAs use natural selection as the inspiration for finding an optimal solution in a near infinite search space. There are several steps to the process described below, and shown in

[1] While the optimization proposed in this section uses a variant on an EA, many other metaheuristic methods such as simulated annealing or hill climbing could have been used. See section A in the appendix for more information regarding EAs.

figure 5.1.

1. **Initialize** a random population of potential solutions.

2. **Evaluate** each solution in the population and assign each an *objective score*.

3. **Select** a subset of the population, usually those which are assigned the "best" score in the previous step.

4. **Vary** the selected subset using Darwinian concepts such as combining and/or mutating the best solutions.

5. **Replace** the previous population with the new, varied population. Then either terminate the process, or take the new population and return to step 2.

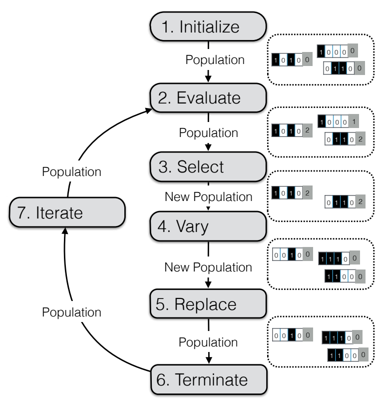

Figure 5.1: Overview of an Evolutionary Algorithm flow

For the purposes of this thesis, the population of solutions refers to either *(a)* transaction sequences, or *(b) audit score sheets*. When evolving transaction sequences, a single *audit score sheet* is selected to **evaluate** the *audit score* for each potential solution. Conversely, a single transaction sequence is selected to **evaluate** each *audit score sheet*against when evolving auditing procedures. Thus, each potential solution is assigned both a measure of taxable income, as well as an *audit score*. These metrics are used to **select** a subset of the best performing solutions, which is then **varied** using methods described further in section 5.4. Finally, the new population **replaces** the previous population, at which point the process is either repeated or terminated.

This process described a "uni-directional" EA because both are being compared to a static objective. That is, all transaction sequences generate their *audit score* based on the same *audit score sheet*. Similarly, all *audit score sheets* are evaluated against the same transaction sequence.

5.3.1 Co-Evolution

Given the two well-defined uni-directional EAs, a "bi-directional" EA, or a co-evolutionary algorithm, can be established. As shown in figure 5.2, a population of individual tax evaders (transaction sequences) is presented with a population of individual auditors (*audit score sheets*). Rather than **evaluating** each solution based on a single objective, multiple solutions from the opposing population are used.

Specifically, the **evaluation** process is expanded by

(i) selecting a size k subset of the opposing population, then

(ii) calculating the total *objective score* as a function of the *objective scores* generated from all k **evaluations**.

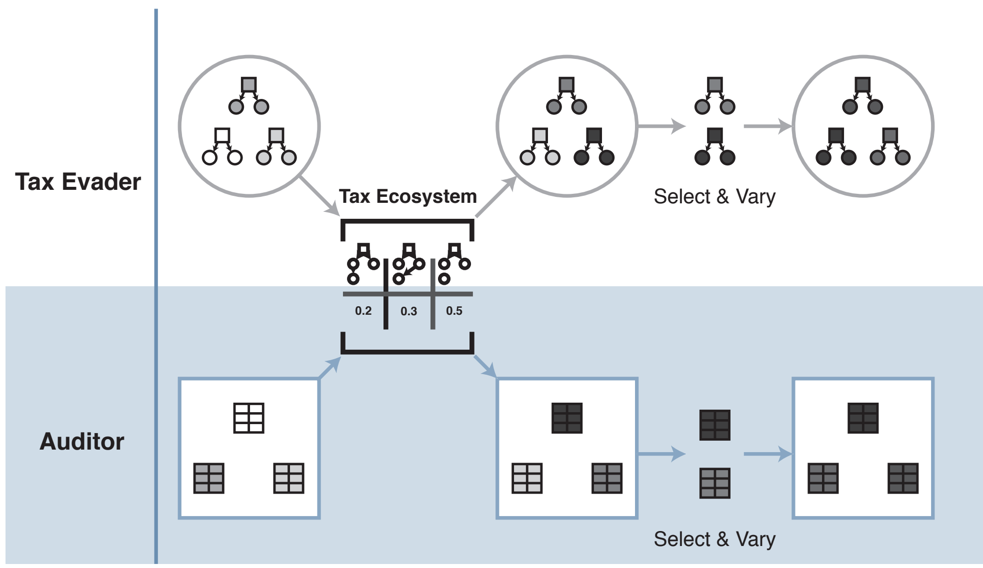

Figure 5.2: A pool of transaction sequence co-evolves with a pool of *audit score sheets* by selecting subsets of one another's populations

In this way, the dynamics between potentially abusive tax shelters and the auditors' response can be studied. This is a step towards the goal of anticipating new iterations of abusive tax shelters mentioned in section 1.2.1.

5.4 Objective Function Optimization

In order to optimize using grammatical evolution as described in appendix A, we must establish a mapping between the set of positive integers and the object that we are attempting to optimize. The following optimization was implemented in Java as both a uni-directional and co-evolutionary search.

We can describe the process by which sequences of transactions and initial ownership network are generated by defining a mapping $\Xi_t : \mathbb{Z}_+^n \mapsto \mathbf{T} \times \Gamma$ that maps a list of n integers to an element in the set of sequences of transaction ($\mathbf{T}$) and an element in the set of all ownership networks (Γ). Thus, for any $\mathbf{x} \in \mathbb{Z}_+^n$, $\Xi_t(\mathbf{x}) = (\mathbf{t}, \gamma_0)$ where $\mathbf{t} \in \mathbf{T}$ is a sequence of transactions and $\gamma_0 \in \Gamma$ is an initial network.

We can now define the space of auditing observables as Ψ, where for some $m \in \mathbb{Z}_+$,

$$\Psi = \{\{b_i\}_{i=0}^{m} : b_i \in [0,1] \text{ and } \sum_{i=0}^{m} b_i = 1\} \subset \mathbb{R}_+^m$$

The map $\Xi_a : \mathbb{Z}_+^m \mapsto \Psi$ maps a vector $\mathbf{y} \in \mathbb{Z}_+^m$ to an element in the set of auditing behavior.

The function $\mathbf{F}$ can be broken up into a network of transition functions that has the same length as the number of transactions in the transaction set contained within the function call (k). Each transition function generates a new network state and an audit score. So for all $i \in [0, k]$, $F_i\left(t_i, \gamma_i, \psi\right) = \left(\gamma_{i+1}, s_i\right) \quad \text{where} \quad s = s_k$

The goal of the tax evader is to minimize both audit likelihood and taxable income. First of all, each set of transactions generates a *taxable income*, η. Secondly an audit score sheet generates an *audit score*, s based on a network of transactions, which represents the likelihood that a scheme will be audited, i.e. the risk of being audited. Thus, we can represent the fitness function, h_e for a tax evasion scheme, given a specific audit score sheet, as $h_e(\eta, s)$

The goal of the auditor is to maximize the likelihood of an audit of a network of transactions with low taxable income. The fitness function for an audit score sheet given a specific tax evasion scheme is a function which reflects such a relationship, represented by $h_a(\eta, s)$. An audit score sheet is fit for a specific evasion scheme if either *1)* there is *not* a suspiciously low amount of taxable income *2)* if there is a high likelihood that if not much tax is collected, then the scheme will be audited

We describe how to judge the fitness of a network of transactions $\mathbf{t}$ and an auditing behavior ψ based on the taxable income η and audit score s generated from the tax ecosystem model $\mathbf{F}$. Now it is possible to fully define the maximizing objectives of networks of transactions as

$$\begin{gathered}\arg\max_{\mathbf{x}^* \in X} \left[h_e\left(\mathbf{F}\left(\Xi_t(\mathbf{x}^*), \Xi_a(\mathbf{y})\right)\right)\right] = \\ \arg\max_{\mathbf{t}^* \in \mathbf{T}, \gamma_0^* \in \Gamma} \left[h_e\left(\mathbf{F}\left(\mathbf{t}^*, \gamma_0^*, \psi\right)\right)\right]\end{gathered}$$

over all $\mathbf{y} \in B(\hat{\mathbf{y}}, r_1)$ for some $\hat{\mathbf{y}} \in \mathbb{Z}_+^m$, where $B(\hat{\mathbf{y}}, r_1)$ is a *ball* of radius $r_1 \in \mathbb{R}_+$ around

$\hat{\mathbf{y}}$. This represents the fact that the goal of the GA is to find local maxima around some subset of auditing behavior, rather than attempting to search the entire Φ space. Conversely, the objective for the auditing behaviors is to maximize the h_a function, i.e. the goal is

$$\arg\max_{\mathbf{y}^* \in \mathbb{Z}_+^m} \left[h_a \left(\mathbf{F}\left(\Xi_t(\mathbf{x}), \Xi_a(\mathbf{y}^*)\right)\right)\right] = \arg\max_{\psi^* \in \Psi} \left[h_a \left(\mathbf{F}\left(\mathbf{t}, \gamma_0, \psi^*\right)\right)\right]$$

over all $\mathbf{x} \in B(\hat{\mathbf{x}}, r_2)$ for some $\hat{\mathbf{x}} \in \hat{X}$, where $B(\hat{\mathbf{x}}, r_2)$ is a *ball* of radius $r_2 \in \mathbb{R}_+$ around $\hat{\mathbf{x}}$. Similar to the previous objective function, this represents the fact that the EA only searches for local maxima around a subset of all transaction sets and initial model states.

6

Demonstration

The purpose of this chapter is to show the different ways in which the representation, simulation and optimization established in Chapters 3, 4 and 5 can be applied to the analysis of an abusive tax shelter known as Distressed Asset Debt, or DAD. The first section of this chapter will describe the tax shelter, along with the government's responses to disallow similar abusive activity [1].

The following analysis is three-fold. First, notation from the formalism described in section 3.1 will show how the flow of taxable income in DAD is captured by the representation. Next, *audit score sheets* will be applied against the DAD transaction sequence to illustrate how the *audit score* changes with different *audit score sheets*. The analysis will be completed with a description of how the optimal *audit score sheet* that fully characterizes DAD can be identified by simulation, as well as a transaction sequence that circumvents the optimal *audit score sheet* by simulation and optimization.

6.1 Distressed Asset Debt (DAD)

The purpose of the DAD scheme was to import highly depreciated assets from abroad without changing their adjusted basis, and sell them to individual taxpayers. This way, the taxpayers could claim deductions on the depreciation in value as if they themselves had incurred the loss, offsetting real gains elsewhere in their portfolio.

As discussed in section 2.1.1, the adjusted basis of an asset *carries over* when contributed to a partnership, thus the promoter of the shelter used this fact to convert a foreign business's worthless assets into tax savings for US investors [2].

[1] As mentioned in the previous three chapters, there exists a preliminary software implementation of the methods

[2] The exact names of the players will be changed, but the mechanics are the same.

Example 6. Parua was at one time the largest retailer of household appliances and consumer electronics in Brazil [7], a position which was partially obtained by offering a substantial consumer credit program. Many customers that purchased goods on credit were unable to pay, thus the company accumulated roughly $30 million in delinquent accounts, otherwise known as *trade receivables*.

Seeking to salvage some profit from these worthless accounts, a partnership Samarth, run by an American tax shelter promoter Roberts, offered Parua a 99% share in the partnership in exchange for contributing the accounts to the partnership (figure 6.1a).

Concurrently, Roberts had established 30 additional TradeCo partnerships [3], each of which Samarth obtained a 99% share of by contributing some of the receivables it obtained from Parua in exchange for a supermajority share, as shown in figure 6.1b. None of the trading companies made a §754 election.

Samarth subsequently contributed its supermajority interest in the trading companies to another set of HoldCo partnership in exchange for a 99% interest in each, shown in figure 6.1c. Because contributions are non-taxable events, the outside basis in each of interests reflected the price of the receivables *without* taking into account their delinquency.

Individual US investors seeking to lower their tax liability were found, who purchased shares in the holding companies (figure 6.1d). The investors subsequently contributed recourse notes to the recently purchased shares in order to increase their outside basis (figure 6.1e) [4]. This did not trigger a §743(b) basis adjustment because Samarth had deliberately *not* made a §754 election. Shortly thereafter, the trading companies claimed partially worthless debt deductions on the receivables that they held from Parua. Those deductions flowed through to the individual investors, who used the loss to offset income elsewhere in their portfolio, shown in figure 6.1f. Samarth claimed losses on the sale of interests in the holding company that it sold to the investors.

In total, roughly $60 million in deductions were claimed from a set of assets which shouldn't have produced any tax benefits in the US.

6.1.1 Government Response

Partially in response to the scheme described in example 6, the US government made two important amendments to the IRC in 2004.

No Transfer of Pre-Contribution Loss: §704(c)(1)(C) was added to the Code, which specifies that if an asset is contributed to a partnership that experienced any decline in value previous to the contribution event, then that loss can only be taken into account by the contributing partner.

This disallowed any DAD-like transactions because a key component is that the contributing partner (Parua) could transfer their loss to other parties (US investors). The aforementioned section disallowed this key component.

[3] In actuality, there were 14 trading companies, but 30 was chosen to simplify the computations

[4] The investor can only claim deductions as large as their outside basis, thus contributing recourse notes allowed them to artificially increase their outside basis to allow deductions to flow through.

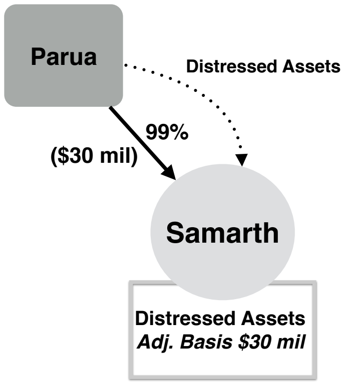

(a) Parua contributes distressed assets

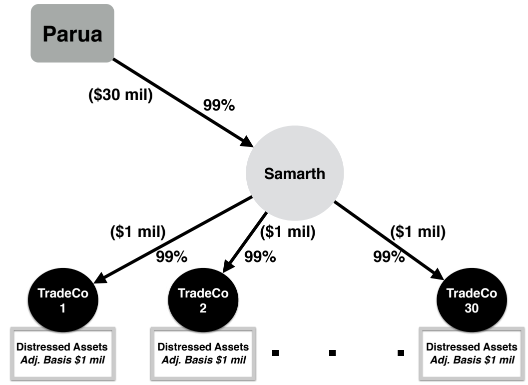

(b) Samarth contributes assets to TradeCos

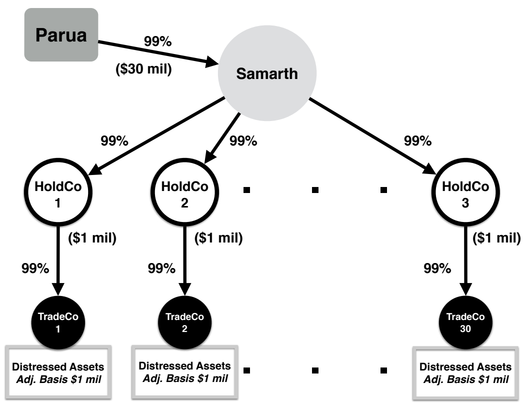

(c) Samarth contributes shares to HoldCos

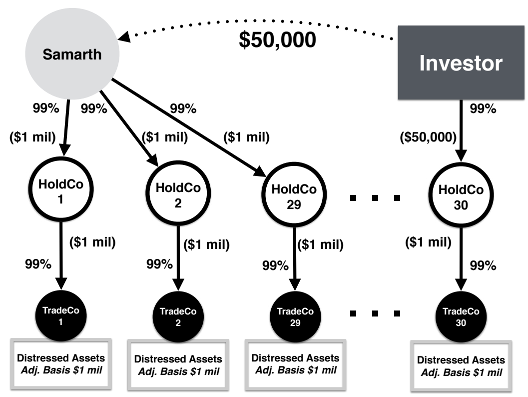

(d) Investors purchase HoldCo shares

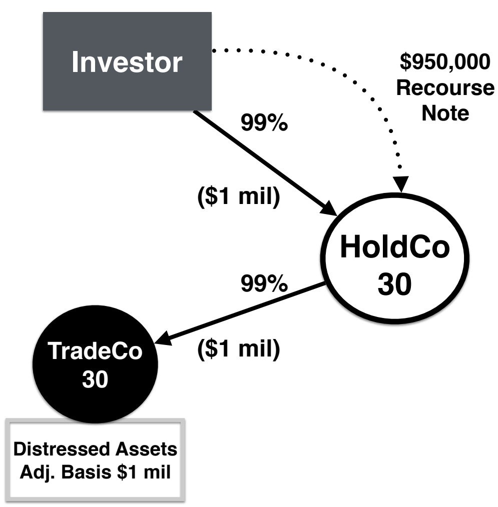

(e) Investor contributes recourse note to increase outside basis

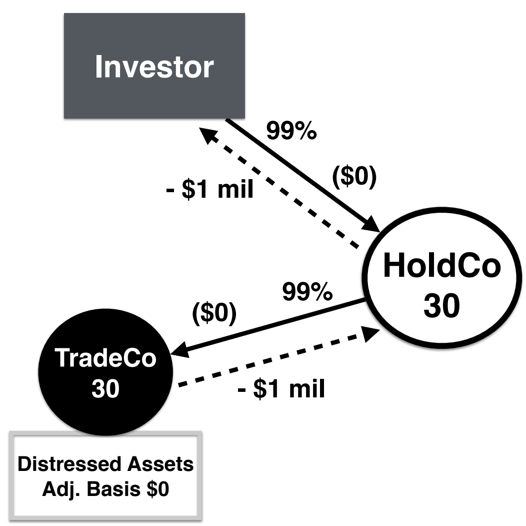

(f) Debt deductions flow to investor

Figure 6.1: Figures that show how the DAD scheme is implemented. Solid lines denote ownership, dotted lines denote an asset transfer and dashed lines denote deductions.

Mandatory §743(b) Basis Adjustment: Several additions to §743 were made that specified conditions under which a partnership is required to adjust the basis of their assets even if they had not previously made a §754 election. If a partnership has a "substantial built-in loss" immediately after the sale of a partnership interest, defined as a loss exceeding $250,000, then the partnership must complete a §743(b) basis adjustment. In a different way, this disallowed DAD by confronting the ability that a partnership had to duplicate losses [5].

In the DAD scheme, Samarth's sale of its interest in the HoldCos would have constituted a substantial built-in loss because its adjusted basis in the partnership interest (Samarth's outside basis of $1 million) exceeded the FMV of the interest ($\sim$ $0 after the partially worthless debt deductions are claimed) by more that $250,000.

6.2 Analysis of DAD

Before conducting the analysis, the transaction sequence that composes DAD must be established. Without loss of generality, it is assumed that there is only one holding company (HoldCo) and one trading company (TradeCo), and that the distressed assets have an adjusted basis of $1 million. In order of occurrence:

1. Parua *contributes* distressed assets with an adjusted basis of $1 million to Samarth in exchange for a 99% share (figure 6.1a).

2. Samarth *contributes* the distressed assets to a partnership TradeCo in exchange for a 99% share (figure 6.1b).

3. Samarth *contributes* its share in TradeCo to HoldCo in exchange for a 99% share in HoldCo (figure 6.1c).

4. An investor *purchases* Samarth's share in HoldCo for $50,000 (figure 6.1d).

5. TradeCo *claims* partially worthless debt deductions, which flow to the investor (figure 6.1f).

6.2.1 Tax Calculation

This section focuses on the acquisition of Samarth's interest in HoldCo, and the subsequent sale to an individual US taxpayer. Recall that before acquiring an interest in HolcCo, Samarth purchased a 99% share in TradeCo by contributing the distressed assets, which have an adjusted basis of roughly $1 million. Thus, their outside basis in their share of

[5]Samarth was able to claim a loss on the sale of its interest in HoldCo, then the investor claimed losses on essentially the exact same assets

TradeCo is \$1 million due to carry-over basis. These characteristics are described below in the established notation described in section 3.1. This serves as a proof of concept to show that the established notation can be applied to real life scenarios.

> TradeCo's only asset (the distressed assets) are represented by the tuple $(b, \tau) = (1mil, 1)$, where $1,000,000$ is its adjusted basis and 1 is the asset type [6].
>
> Samarth's interest in TradeCo is described by the tuple $(\theta, \kappa, s, \sigma) = (0, 0, 0.99, 1mil)$. Currently, Samarth has generated no income of any kind, has a share of 99% in TradeCo and has an outside basis of \$1 million.

Contributing a partnership interest is treated identically to the contribution of any other asset, with the outside basis in the partnership acting as the adjusted basis. So when Samarth contributes its share of TradeCo to HoldCo in exchange for a 99% share, the notational changes are minor.

> HoldCo's sole asset is its interest in TradeCo, which is described as $(b, \tau) = (1mil, 1)$ [7].
>
> Samarth no longer has an interest in TradeCo, but its interest in HoldCo is the same as its interest in TradeCo, namely $(\theta, \kappa, s, \sigma) = (0, 0, 0.99, 1mil)$.

Samarth then sells its interest in HoldCo to a US investor for $\$50,000$ [8], significantly less that its adjusted basis of \$1 million.

> The total taxable income that is pushed to Samarth is the difference between the FMV of its interest and its outside basis in HoldCo
>
> $$I_t = \bar{e}_s \sum_{a \in A} F(a,t) - \bar{e}_\sigma = 0.99 * (50,000) - 1,000,000 = -950,500$$
>
> The amount of which is ordinary is calculated by summing over the built-in income and accrued income of all ordinary assets within the partnership
>
> $$I_\theta = \sum_{a \in A_0 \cup A_1} \Phi(\bar{e}, a) + \Gamma(\bar{e}, a) = -950,000$$
>
> Finally, the capital income is defined as the difference between the total taxable income and the ordinary income, i.e.
>
> $$I_\kappa = I_t - I_\theta = -950,500 + 900,000 = -50500$$

[6] distressed assets are trade receivables, and are thus treated as ordinary assets as dictated by §751

[7] the asset type remains 1 because the underlying asset that generated the interest, namely the distressed assets, are of type 1

[8] Presumably, Samarth claims that the FMV of the distressed assets have substantially declined, but have not reached zero yet

Observable	Independent	Random	Joint
Pship Interest Sale (1)	0.33	0	0
No §754 Election (2)	0.33	0.15	0
Substantial Built-in Loss (3)	0.33	0	0
(1) and (2)	0	0	0
(1) and (3)	0	0.5	0
(2) and (3)	0	0.35	0
(1) and (2) and (3)	0	0	1.0

Table 6.1: Three sample *audit score sheets*

> Samarth is now described by the tuple $(\theta, \kappa, s, \sigma) = (-950,000, -500, 0.0, 0.0)$ in respect to HoldCo. The US taxpayer is now described, in respect to HoldCo, as the tuple $(\theta, \kappa, s, \sigma) = (0, 0, 0.99, 50,000)$.

Thus, Samarth's ability to claim a large loss from the sale of the interest is shown from the representation described in section 3.1.

6.2.2 Audit Score Calculation

The goal of this section is to show how certain *audit score sheets* are more optimal than others. This is accomplished by comparing three different *audit score sheets* against the sequence of DAD transactions, as well as a sequence of sample non-abusive transactions, to illustrate how the *audit score* calculations take place.

Recall the *audit score sheet* used as an example in section 4.3. The three *observables* used were:

(1) Sale of a partnership interest

(2) Transfer of a partnership interest which has not made a §754 election.

(3) There exists a substantial built-in loss through a partnership interest immediately before the transaction occurs.

Thus an *audit score sheet* corresponding to these three *observables* has seven rows, as shown in table 4.1, that represent all joint combinations of events [9].

These *observables* were chosen because together, they characterize a key component of DAD. Specifically, the sale of a partnership asset, from a partnership that has *not* made a §754 election, that results in a substantial built-in loss. Each *observable* can occur independently, but signify abuse when combined.

Consider the three hypothetical *audit score sheets* shown in table 6.1.

[9] As mentioned in section 4.3, an *audit score sheet* representing n individual *observables* will have $2^n - 1$ rows.

All three *audit score sheets* will generate an *audit score* of one when evaluated against the sequence of DAD transaction because *observables* (1), (2) and (3) all occur in the same transaction.

Independent Sums up the *audit points* associated with each event happening individually.

Random Sums up *audit points*of (2) occurring independently, (1) and (2) occurring together, and (1) and (3) occurring together.

Joint Sums the single *audit point*of all three events occurring in the same transaction

Now consider a routine transaction which should not raise any suspicion from auditors. Specifically, suppose that HoldCo *had* made a §754 election, which would not generate any artificial deductions for the investors. In this case, *observables* (1) and (3) occur, but not (2). The **independent** *audit score sheet* will generate an *audit score* of $0.33 + 0.33 = 0.66$, the **random** *audit score sheet* will generate an *audit score* of 0.5, and the **joint** will generate a zero valued audit score.

It is clear that the **joint** *audit score sheet* is the best option because it generates a high *audit score* when a transaction is suspicious and a low *audit score*otherwise. While the **independent** and **random** *audit score sheets* also detect the DAD scheme, they cannot distinguish routine transactions that contain certain characteristics also associated with DAD. Since IRS auditors face significant resource constraints [15], avoiding false positives is of great concern. Thus the goal of the *audit score sheet* is to characterize abusive transactions in the most compact manner possible.

In reality, abusive transactions are rarely able to be characterized perfectly by a joint combination of *observables*. Many elements that compose such transactions cannot be easily observed, and thus cannot fit neatly into an *audit score sheet*. Still, an *audit score sheet* can be effective by characterizing an abusive tax shelter by the most accurate joint occurrence of *observables*. This will likely result in false positive detections on routine transactions, but can still be useful for constructing characterizations of abusive tax shelters.

6.2.3 Finding Optimal Solutions

This section will show that optimal solutions exist within the search space, whether that be transaction sequences or *audit score sheets*. This will be accomplished by first describing an *audit score sheet* that characterizes the DAD transaction sequence in an optimal manner. Next, a transaction sequence that exists in the search space will be shown, which accomplishes the same tax consequences as the original DAD scheme, but evades the sample *audit score sheets* shown in table 6.1.

6.2.3.1 Optimal *audit score sheet*

Consider the example *audit score sheets* in table 6.1, as well as the two scenarios below that describe a series of transactions similar to DAD.

Transaction Sequence (*i*) HoldCo makes a §754 election upon formation. Thus, the adjusted basis of the distressed assets are adjusted downward upon the investor's purchase of a 99% share in HoldCo. This nullifies the deductions for the investor, leaves them with 0 in taxable income and constitutes a non-abusive transaction.

Transaction Sequence (*ii*) HoldCo does *not* make a §754 election, which generates a deduction for the investor of $\$1,000,000$ (or $-\$1,000,000$ in taxable income). This is an abusive transaction that should generate a high *audit score*.

Suppose further that the US investor has $5 million in taxable income. Thus, if they engage in transaction sequence (*ii*), then they will only have $\$5,000,000 - \$1,000,000 = \$4,000,000$ in taxable income, whereas transaction sequence (*i*) results in no tax savings.

Recall from section 5.4 that the objective function for the *audit score sheet* is some $h_a(\eta, s)$, where η is the taxable income of the investor and s is the *audit score* associated with the *audit score sheet*. The objective function should be chosen such that:

(a) when compared against the same *abusive* transaction sequence, *audit score sheets* that generate a high *audit score* should be assigned a high objective value, and

(b) when compared against the same *non-abusive* transaction sequence, *audit score sheets* that generate a low *audit score* should be assigned a high objective value.

Thus, an *audit score sheet* objective function of $h_a(\eta, s) = \eta(1 - s)$ will be used in this scenario. Table 6.2 below displays the h_a objective value for each combination [10].

		Transaction Sequence	
		(*i*)	(*ii*)
Audit score sheet	**Independent**	1,700,000	0
	Random	2,500,000	0
	Joint	5,000,000	0

Table 6.2: *Audit score sheet* objective function values from transaction sequence and *audit score sheets* combinations

The objective values for the *audit score sheets* are all the same when compared against transaction sequence (*ii*) because, as shown in section 6.2.2, each one is able to capture the three *observables* that describe DAD.

More important are the objective values associated with transaction sequence (*i*). The **joint** *audit score sheet* has the highest objective value in that scenario because it assigns a

[10] The *audit scores*from each *audit score sheet*/transaction sequence combination is found in section 6.2.2

low *audit score* (0) to a routine transaction sequence. The other two *audit score sheets*, on the other hand, assign a higher *audit score* (0.66 and 0.5 respectively), resulting in a lower objective value.

While the optimality of the **joint** *audit score sheet* was shown intuitively in section 6.2.2, this section has shown that the proper objective function can determine the optimal *audit score sheet* as well. This fact establishes the foundation for a computational evaluation of a near infinite search space.

6.2.3.2 Optimal Transaction Sequence

Here, a transaction sequence will be presented which generates \$1 million of deductions for the US investor, but in a manner which is not detectable by the optimal *audit score sheet* discussed in section 6.2.2. Given that the transaction sequence exists within the integer search space, a directed search will be able to find the optimal transaction sequence that evades the *audit score sheets*.

Consider the following variation on the DAD scheme described in section 6.1.

> After receiving the distressed assets with an adjusted basis of \$1 million, Samarth splits the assets into five pieces, each of which have an adjusted basis of $\$200,000$. Samarth then contributes each piece to a different partnership TradeCo $1-5$ in exchange for a 99% share. Then, in exchange for the supermajority interest in TradeCo i ($i \in [1,5]$), partnership HoldCo i gives Samarth a 99% interest. Finally, the investor purchases Samarth's share in all of the TradeCo partnership.

Note that this is nearly identical to the DAD scheme considered in the previous section, except the distressed assets are split into smaller pieces, then separated into different partnerships. Because each piece has an adjusted basis of less than $\$250,000$, the investor's purchase of each individual TradeCo interest does not trigger a mandatory basis adjustment [11].

Thus, the optimal *audit score sheet* found in section 6.2.3.1 is not able to detect this minor variation on DAD. Because *observable* (3) is not present, the final transaction is only characterized by the sale of an interest in partnership that has not made a §754 election. Because this is a routine transaction that rarely contains abuse, the *audit score sheet* is unable to effectively characterize the new variation.

The existence of this variation is significant because the sequence of transactions that compose the scheme is able to be represented by the grammar described in appendix A. Thus, it exists in the search space defined in section 5.4, and can be found by a properly configured search heuristic. Furthermore, the evolved transaction sequence illustrates how

[11] Due to the 2004 of §704$(c)(1)(C)$, this variation would still be disallowed, but not due to the mandatory basis adjustment rules

audit score sheets can be evaded if they do not contain all of the relevant *observables*. That is, the *audit score sheets* will be unable to fully characterize certain schemes.

In the DAD example, the new optimal transaction sequence involves splitting up the distressed assets such that each individual partnership generates a $200,000 deduction, less than the $250,00 threshold that is detectable. Thus, without decreasing the threshold for a substantial built-in loss [12] , no *audit score sheet* can properly characterize the new scheme with the current three *observables*.

[12] Decreasing the threshold poses an interesting problem. $250,000 was chosen presumably because it seems unlikely that a loss that large could occur without being engineered. But decreasing the threshold too low could cause unnecessary logistical costs for partnerships which encounter such a loss under routine circumstances

7

Implications

While this method was developed originally to anticipate new iterations of abusive tax shelters, the modular design allows for many use cases across different domains. Many benefits of the approach are rooted in the ability to study how the tax attributes of complex partnership structures respond to financial activity. This can add an element of clarity to partnership taxation for the involved parties.

7.1 Benefit for Taxpayers

Non-compliance with the US tax code is not always an intentional act. While many professional tax preparers may knowingly construct transaction sequences that are potentially non-compliant, a significant amount of non-compliance is unintentional. This may be due to untrained tax preparers, a needlessly complicated compliance process, but likely a combination of the two. This thesis presents potential solutions to such issues

Ease of Compliance

> Previous studies have shown that given a choice between complying with the tax code and engaging in potentially abusive activity, individual US taxpayers will choose to comply more than taxpayers in most other countries [6]. This implies that ease of filing can have a significant, positive effect on compliance rates.
>
> Software such as TurboTax or H&R Block, as well as free online services like freefile have allowed individual taxpayers to file tax returns with minimal background knowledge on US's tax law. While many taxpayers continue to hire accountants, particularly those that are self-employed, this software has greatly increased taxpayers' access to filing ability.

Similarly, partnership tax software could afford the same opportunity to small businesses. Accountants and tax professionals would still play a vital role in the filing of K-1 returns, given the many complex partnership taxation laws, but small business owners could explore the tax implications of different types of activity without a deep knowledge of the tax code. The representation of partnership taxation in chapter 3 serves as a model by which the allocation of taxable income and the adjustment of bases could be incorporated into user-friendly software.

For partnership tax professionals, such software could serve as a computational tool that takes care of the "dirty work". The method described in this thesis cannot replace the job of tax preparers, but it can perform many tedious tasks that are unnecessary for a trained professional to engage in. Similar to the way in which Computer-Aided Design (CAD) programs assist engineers, partnership tax software has the potential to allow tax professionals to spend more time on business organization and operation, rather than calculating §743(b) basis adjustments and other logistical tax issues.

Positive Compliance Externalities

Voluntary compliance by US taxpayers is strongly influenced by whether or not those around them are complying. That is, if taxpayers feel as though others are not paying their "fair share", then they will be significantly more inclined to engage in potentially abusive activity. This has been both shown as quantitative model results [8, 24, 17] and stated by the US government as an argument for increased auditing efficiency [18].

Increasing the IRS's ability to effectively audit partnerships can inspire faith in the American people that their tax system is fair, which will contribute to increased voluntary compliance.

7.2 Benefit for Auditors

The Government Accountability Office (GAO) has been greatly concerned with the IRS's ability to effectively understand and audit large partnerships as discussed in section 2.1.2. Facing large budget cuts, political scrutiny and pressure to collect revenue, the IRS must develop new tools in its arsenal to audit large, complex partnership structures.

Many of the challenges faced by the IRS are addressed by the methodology proposed in this thesis.

Increase Electronic Filing

Partnerships with 100 partners or less are generally not required to electronically file (e-file), which limits the IRS's ability to effectively conduct partnership audits [15].

A primary reason for this is the assumption that e-filing is a costly process. But as discussed in section 7.1, the propagation of user-friendly partnership tax software has the capacity to reduce costs of e-filing to smaller partnerships. Not only with this likely result in a greater compliance rate of taxpayers, but it will increase the IRS's access to tax data. This will both increase the auditing rate of non-compliant returns, while reducing the rate at which compliant returns are audited.

More Efficient Audits

The IRS examination, or audit, process is conducted in three separate steps [15]. *Return selection* is a generally automated process in which risky tax returns are selected for further review. During the *classification* process, trained staff review the selected returns to determine which pose a high enough risk of non-compliance to be examined further. The final step, an *examination* of the tax return is conducted, which usually consists of additional contact with the taxpayers as well as a resource intensive information evaluation process.

Given that the *return selection* process is automated, the representation of audit policy presented in chapter 4 can offer additional insight into which returns are selected. Particularly in respect to electronically filed returns, numerical weights or *audit points* can be assigned to information included in the return to indicate potentially non-compliant behavior. Once the *audit score* passes a certain threshold, the return may proceed to the next step.

Unlike the automated *return selection* process, the method proposed in chapter 4 can serve more as a heuristic to trained staff during the classification process. The *audit score* process can highlight areas of concern to such professionals, who can then use their expert knowledge to determine whether the automation uncovered legitimate areas of concern. Not only will a positive detection result in increased efficiency, but the experts can use false detections to improve the *audit points* assigned to certain types of behavior.

The more resource heavy *examination* step can be aided by reducing the amount of information required to conduct an efficient audit. In the words of an IRS auditor [14]:

> ...[Because] income is pushed down so many tiers, you are never able to find out where the real problems or duplication of deductions exist. The reporting of income, expenses could be duplicated but there is no way to figure it out unless you drill down and audit all tiers, all tax returns.

Rather than requiring an audit of every entity involved in a complex partnership structure, the methodology described in chapter 5 can help to classify abusive behavior with a minimal amount of information. Thus, a suspicious structure involving 100 partnerships may only require a small percentage of them to be audited in order to determine areas of potential error or abuse. A reduction of the time to conduct an efficient audit is of great concern to the IRS [13], particularly considering the over 800% increase in partnerships exceeding 500,000 partners in size. Extending the proposed methodology to operate in this manner would require iterating with partnership tax experts to determine the most effective *observables* and will be left for future work.

Learning Tool The complexities of partnership taxation are a main source of the IRS's difficulty, given that many of their staff are not trained to handle such as complicated aspect of the tax code [14]. Developing a learning tool to understand partnership taxation in conjunction with the representation developed in chapter 3 can increase the staff's ability to audit large partnership structures. Substantial prior work has shown that computer modeling can greatly aid the learning process [19, 5, 31]. Abstractly representing a complex system and investigating the effects that certain transactions have on a large partnership may be the most effective way to train IRS auditors.

This would greatly reduce the amount of resources required to conduct audits because the IRS generally requires TEFRA coordinators (arising from the 1982 Tax Equity and Fiscal Responsibility Act) and partnership specialists to evaluate complexities [14]. The IRS would thus be much less likely to run out of time due to restrictive statutes of limitations on certain types of behavior. The development of a sophisticated graphical user interface is required to fill this role.

7.3 Benefit for Policy

Congress, as well as groups such as the New York Bar Association and American Bar Associated (ABA) that advise Congress on tax policy, have an opportunity to strengthen their insight by incorporating computational methodologies. That is, tax minimizing behavior not intended by policy-makers is often implemented through pre-existing rules, many of which were made explicitly to prevent similar behavior that had been observed in the past [35, 21]. The methodology presented in this thesis can be used to prevent future potential abuses of the tax code in a number of ways.

Anticipation of Tax Avoidance

Many abusive tax shelters or other unintended effects of partnership tax policy arise due to technical interactions between clauses that were not considered during the policy making process [21]. These interactions are often discovered by professional tax planners with a deep understanding of the tax code [26, 35]. The methodology presented in chapter 5 is an automated process that models the discovery process that such tax planners engage in. Potential changes can be either be incorporated into the representation from section 3.1, or modeled as changes in audit procedure as described in chapter 4.

Rather than replacing tax experts, the method serves as a supplement to partnership policy specialists by calculating numerical tax implications arising in large partnership structures. The use of highly tiered partnership structures results in often times unknown implications of changes to policy. By using this thesis's methodology, policy-makers could not only anticipate the effect that proposed policy would have, but could use it to construct new policy. Increased sophistication of relevant *observables*, as well as the completion of the representation described in chapter 3 is required to develop this capability.

Reconciliation of Entity and Aggregate Models

As mentioned in chapter 3, the proposed abstract representation of partnership taxation has not been proven to fully reconcile the two models. That being said, there is no evidence that the representation could *not* serve as an alternative or supplement to current accounting models. The basis adjustments described in §734 and §743 arguably arise from the need to equalize entity and aggregate concepts of *outside basis* and taxable income [11]. Combining the models into a single representation could mitigate the confusion surrounding the issue.

Even if the new representation is not used by tax professionals, it could shed light on certain illogical aspects of the current method for basis adjustment. As pointed out by Professor Howard Abrams, there is a fatal flaw in the implementation of §734(*b*) basis adjustments. Unlike §743(*b*), the application of basis adjustments in the case of distributions is not in respect to specific partners [1]. This can cause inequity problems in non-liquidating distributions for technical issues discussed further by Abrams in his paper. While the new representation cannot rectify the problem, it can help policy-makers understand the origins of the inequity and potentially make policy to mitigate them, similar to the changes made to §743(*b*) basis adjustments.

Elimination of the §754 Election

In the words of Stephen Shay, "...elections always result in a loss for the fisc [Treasury]..." [28], where an election in the tax code allows taxpayers to opt in or out of a certain set of rules. If taxpayers are given an option, they will generally take the route that leads to a lower tax liability, decreasing collected revenue. If this is the intended effect of policy, then elections do not pose a problem. This is not the case for §754 elections.

As discussed in section 2.1.1.1, the likely reason for originally making basis adjustments optional was the logistical difficulties that small partnerships would have with the calculation of the adjustments [28]. The software proposed earlier in section 7.1 would eliminate the resources necessary to calculate the adjustments. Eliminating the election would require basis adjustments to take place upon the occurrence of the sale of a partnership interest, distribution, recognition of §731(a) gain or loss and the death of a partner subject to §1014(a). A *de minimus* clause would be included to ensure that sufficiently minor adjustment would not result in a waste of resources. Such clauses are common within US tax law and can prevent many problems associated with relatively minor taxation issues.

8

Conclusions and Future Work

The rise in importance of large partnership structures poses opportunities for those with a deep understanding of partnership taxation to exploit ambiguities in the tax law for their own benefit. Conversely, a significant portion of partnership filers in the US attempt to comply, but are constricted by the complexity of partnership taxation. Disallowing abusive behavior, while clarifying complicated taxation scenarios can generate positive effects for both the US government and its citizens.

This thesis has proposed a method towards achieving this goal by answering the three research questions in section 1.2.1. The first task was to determine the correct allocation of taxable income, gain, loss, deductions are credits arising from transactions occurring within partnership structures. Chapter 3 established a formal representation of partnership taxation which can serves as a functioning tax calculator. Next, this thesis aimed to simulate the process of an IRS auditor assigning risk to a sequence of transactions and a partnership structure. To accomplish this, a numerical representation of the auditing procedure, as well as a notion of a tax network, was described in chapter 4.

Finally, the goal of anticipating new forms of abusive tax behavior was proposed. This required a method of optimizing both effective tax minimizing transaction sequences, as well as IRS auditing policies. A search space and objective function for both types of solutions was established in chapter 5, which establishes a process by which a metaheuristic can find optimal transaction sequences and auditing procedures, a step towards anticipation of new abusive tax shelters. The three aspects of the methodology were combined in chapter 6 in order to show the use-cases of potential software built around the approach in this thesis.

Potential implications of the proposed methodology for taxpayers, auditors and policy-makers were then discussed in chapter 7. A standard method for calculating tax presents an opportunity to simplify the process of filing for tax preparers of large partnerships. Increased filing ease, as well as a numerical representation of audit risk, allows IRS auditors to supplement expert knowledge with an automated process. Finally, simpler partnership tax filing, combined with improved analysis tools, allows policy-makers to both better explore policy options, as well as update vestigial elements of Subchapter K.

Each aspect of the presented methodology contains opportunities for extension and improvement. The representation in chapter 3 does not contain the calculation of the effect that liabilities or depreciation deductions have on taxable income. While the notation was constructed with these extensions in mind they have yet to be implemented. Additionally, a small number IRC clauses such as §737 have yet to be built into the representation.

In order to fully encapsulate the aspects of partnership transactions which are both detectable by auditors and potentially indicative of abuse, further discussions with partnership taxation experts is necessary. The *observables* used in this thesis are useful for conceptualizing the approach, but a more realistic list of events must be composed in order to fully realize the potential benefits outlined in chapter 7. Furthermore, additional work should be explored on advanced Artificial Intelligence methods to evolve *observables* themselves through unsupervised learning methods.

As discussed in section 2.2, the characterization of illegal tax behavior can be a highly non-trivial task. Further, many aspects of Subchapter K allow for seemingly inequitable results due to the complex nature of flow-through tax entities. Thus, this thesis is as much about the problems that arise in US partnership taxation law as it is about those who exploit them for financial gain.

Appendices

A

Genetic Algorithms

Co-evolutionary search heuristics [9, 12, 34, 30] evaluate an individual solution based on interactions between populations of multiple solutions. The individual solution may appear good in one context and poor in another context, e.g. one solution's ranking in a population can change depending on other solutions.

A.1 Grammatical Evolution

Grammatical Evolution (GE) is a version of the Genetic Algorithm with a variable length integer representation and a compressed form of indirect mapping using a grammar [23]. We can map to a transaction sequence by means of a grammar which conveniently expresses all possible transaction sequences compactly. GE has an explicit mapping step (genotype-to-phenotype) and biases the search by changing the grammar, e.g. alter the search space size and reduce source code modification. The grammar rewrites the input (genotype) to the output (phenotype), as shown in Figure A.1. Recursive rules in the grammar indicate that the search space (language) is bounded only by the length of the input (genome) used in rewriting.

In GE, the compressed form of the search space is represented by a Backus-Naur Form (BNF) grammar which defines the language that describes the possible output sentences. a BNF grammar has terminal symbols, non-terminal symbols, a start symbol and production rules for rewriting non-terminal symbols. The grammar is used in a generative approach and the production rules are applied to each non-terminal, beginning with the start symbol, until a complete program is formed. The list of integers (genotype) rewrites the start symbol into a sentence. An integer from the list of integers is used to choose a production rule from the current non-terminal symbol by taking the current integer input and the modulo of the current number of production choices. Each time a production from a rule with more than one production choice is selected to rewrite a non-terminal, the next integer is read and the

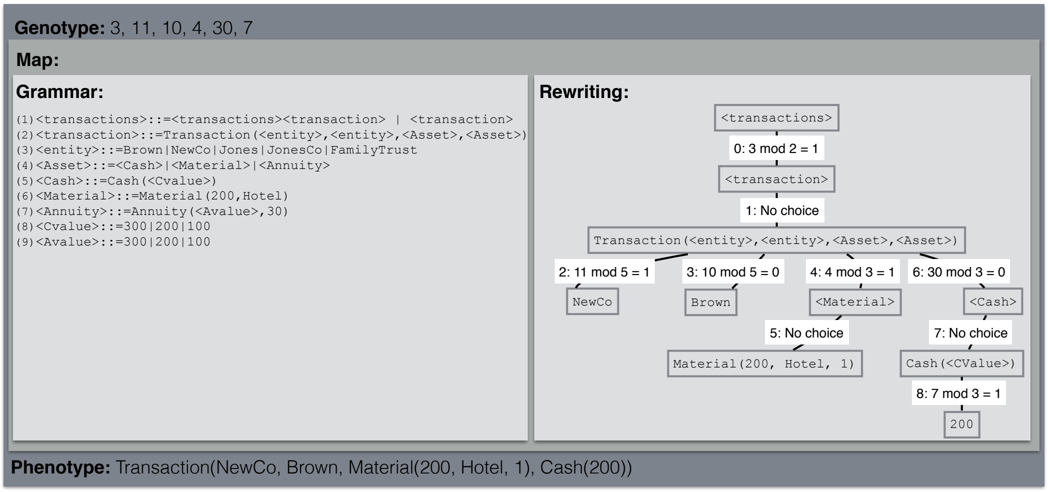

Figure A.1: Example of mapping a list of integers (Genotype) into a list of transactions (Phenotype) by using grammatical evolution

system traverses the genome. The rewriting is complete when the sentence comprises only terminal symbols.

In Figure A.1 there is an example of the rewriting of an integer list (genotype) to a sentence (phenotype) describing a transaction between two entities that exchange assets.

1. We pick the first rule in the grammar as the start symbol, in this case `(1) <transactions>`.
2. Expand the left most non-terminal symbol in our sentence `<transactions>`. We take the current integer input 3 and the modulo of the number of production choices 2, which is 1, thus we pick `<transaction>` the production choice at position 1 (the indexing starts at 0) and rewrite the `<transactions>` with `<transaction>`.
3. Again expand the left most non-terminal symbol `<transaction>`. There is only one production choice here, so it is rewritten to `Transaction(<entity>, <entity>, <Asset>, <Asset>)`.
4. Again expand the left most non-terminal symbol `<entity>`. We take the current integer input 11 and the modulo of the number of production choices 5, which is 1, thus we pick `NewCo`. The sentence is now `Transaction(NewCo, <entity>, <Asset>, <Asset>)`.
5. The left most non-terminal symbol is again `<entity>`. We take the current integer input 10 and the modulo of the number of production choices 5, which is 0, thus we pick `Brown`. The sentence is now `Transaction(NewCo, Brown, <Asset>, <Asset>)`.
6. The left most non-terminal symbol is now `<Asset>`. We take the current integer input 4 and the modulo of the number of production choices 3, which is 1, thus we pick `<Material>`. The sentence is now `Transaction(NewCo, Brown, <Material>, <Asset>)`.
7. The left most non-terminal symbol is now `<Material>`. There are no choices for `<Material>` so we rewrite it with `Material(200, Hotel, 1)`. The sentence is now `Transaction(NewCo, Brown, Material(200, Hotel, 1), <Asset>)`.

8. The left most non-terminal symbol is again `<Asset>`. We take the current integer input 30 and the modulo of the number of production choices 3, which is 0, thus we pick `<Cash>`. The sentence is now `Transaction(NewCo, Brown, Material(200, Hotel, 1), <Cash>)`.
9. The left most non-terminal symbol is now `<Material>`. There are no choices for `<Cash>` so we rewrite it with `Cash(<Cvalue>)`. The sentence is now `Transaction(NewCo, Brown, Material(200, Hotel, 1), Cash(<CValue>)`.
10. The left most non-terminal symbol is `Cash(<CValue>`. We take the current integer input 7 and the modulo of the number of production choices 3, which is 1, thus we pick 200. The sentence is now `Transaction(NewCo, Brown, Material(200, Hotel, 1), Cash(200)`. x
11. There are no more non-terminal symbols left to rewrite and our string rewriting is done.

B

Subchapter K: Relevant Sections

Sale of Partnership Interest: §743(b) and (c)

(b) Adjustment to basis of partnership property

In the case of a transfer of an interest in a partnership by sale or exchange or upon the death of a partner, a partnership with respect to which the election provided in section 754 is in effect or which has a substantial built-in loss immediately after such transfer shall

(1) increase the adjusted basis of the partnership property by the excess of the basis to the transferee partner of his interest in the partnership over his proportionate share of the adjusted basis of the partnership property, or

(2) decrease the adjusted basis of the partnership property by the excess of the transferee partners proportionate share of the adjusted basis of the partnership property over the basis of his interest in the partnership.

Under regulations prescribed by the Secretary, such increase or decrease shall constitute an adjustment to the basis of partnership property with respect to the transferee partner only. A partners proportionate share of the adjusted basis of partnership property shall be determined in accordance with his interest in partnership capital and, in the case of property contributed to the partnership by a partner, section 704 (c) (relating to contributed property) shall apply in determining such share. In the case of an adjustment under this subsection to the basis of partnership property subject to depletion, any depletion allowable shall be determined separately for the transferee partner with respect to his interest in such property.

(c) Allocation of basis

The allocation of basis among partnership properties where subsection (b) is applicable shall be made in accordance with the rules provided in section 755.

Distribution: §734 (b) and (c)

(b) Method of adjustment

In the case of a distribution of property to a partner by a partnership with respect to which the election provided in section 754 is in effect or with respect to which there is a substantial basis reduction, the partnership shall

(1) increase the adjusted basis of partnership property by

- **(A)** the amount of any gain recognized to the distributee partner with respect to such distribution under section 731 (a)(1), and
- **(B)** in the case of distributed property to which section 732 (a)(2) or (b) applies, the excess of the adjusted basis of the distributed property to the partnership immediately before the distribution (as adjusted by section 732 (d)) over the basis of the distributed property to the distributee, as determined under section 732, or

(2) decrease the adjusted basis of partnership property by

- **(A)** the amount of any loss recognized to the distributee partner with respect to such distribution under section 731 (a)(2), and
- **(B)** in the case of distributed property to which section 732 (b) applies, the excess of the basis of the distributed property to the distributee, as determined under section 732, over the adjusted basis of the distributed property to the partnership immediately before such distribution (as adjusted by section 732 (d)).

Paragraph (1)(B) shall not apply to any distributed property which is an interest in another partnership with respect to which the election provided in section 754 is not in effect.

(c) Allocation of basis

The allocation of basis among partnership properties where subsection (b) is applicable shall be made in accordance with the rules provided in section 755.

Rules for allocation of basis: §755(a)

(a) General rule

Any increase or decrease in the adjusted basis of partnership property under section 734 (b) (relating to the optional adjustment to the basis of undistributed partnership property) or section 743 (b) (relating to the optional adjustment to the basis of partnership property in the case of a transfer of an interest in a partnership) shall, except as provided in subsection (b), be allocated

(1) in a manner which has the effect of reducing the difference between the fair market value and the adjusted basis of partnership properties, or

(2) in any other manner permitted by regulations prescribed by the Secretary.

C

Entity vs. Aggregate Model

Under a pure entity model, the entity computes its taxable income without reference to any tax attributes of its beneficial owners. That is, the entity's basis in its assets is independent of any asset basis of its shareholders or partners, the entity has a taxable year and method of accounting independent of the taxable years and accounting methods of its owners, and so on. Once the entity's taxable income or loss is computed, it is passed through to its owners.

Under a pure aggregate model, the entity has no tax attributes of its own and does not compute taxable income. Instead, the entity is treated as no more than an aggregate of its owners, so that, for example, it has no independent basis in its assets but instead uses the aggregate of its owners' bases. Thus, a separate depreciate schedule must be maintained for each owner's interest in each entity asset, each owner may have a distinct holding period for each asset, and so on. Although a pure aggregate approach to pass-thru taxation is possible, in practice it would be cumbersome to implement.

Subchapter K constructs a modified entity model for the taxation of partnerships and their partners. Computationally, we treat a partnership as an entity separate from its partners having its own tax attributes including basis, taxable year, and accounting method. However, much of subchapter K is devoted to injecting some of the aggregate model into the mixture.

For example, §§722 − 723 ensure that initial aggregate inside basis equals initial aggregate outside basis. Further, the rules of §§705 and 752 struggle to preserve this equality. Indeed, the optional basis adjustments of §§734(b) and 743(b) can ensure that the equality between aggregate inside and aggregate outside bases is maintained despite excess distributions and dispositions of partnership interests. Equality of aggregate inside and outside bases minimizes the impact of the partnership of the taxation of its partners.

Bibliography

[1] H. Abrams. The section 734(b) basis adjustment needs repair. *The Tax Lawyer*, 57:209–240, 2004.

[2] H. Abrams. Now you see it, now you don't: Exiting a partnership and making gain disappear. *Emory Public Law Research Paper*, 2009.

[3] H. Abrams. Interview with howard abrams, march 6, 2015. Personal Communication, March 2015.

[4] O. Badar. Simulating tax evasion using agent based modelling and evolutionary search, 2014.

[5] P. Bilkstein and U. Wilensky. Less is more: Agent-based simulation as a powerful learning tool in materials science. In *IV International Joint Conference on Autonomous Agents and Multiagent Systems*. AAMAS, 2005.

[6] E. Christie and M. Holzner. What explains tax evasion? an empirical assessment based on european data. Technical Report 40, The Vienna Institute for International Economic Studies, 2006.

[7] U. T. Court. Superior trading, llc, jetstream business limited, tax matters partners, et al, petitioners v. commisoner of internal revenue, respondent, 2011. `http://www.ustaxcourt.gov/InOpHistoric/superiortrad.TC.WPD.pdf`.

[8] J. S. Davis, G. Hecht, and J. D. Perkins. Social behaviors, enforcement, and tax compliance dynamics. *The Accounting Review*, 78(1):39–69, 2003.

[9] E. D. de Jong, K. O. Stanley, and R. P. Wiegand. Introductory tutorial on coevolution. In *GECCO (Companion)*, pages 3133–3157, 2007.

[10] M. Desai. Complexity of the tax code. `http://media.wix.com/ugd/b391a0_e698250b09da4804945d51119ae2010f.pdf`, March 2014. Testimony of Mihir Desai before the US Senate Committee on Finance.

[11] R. Doernberg, H. Abrams, and D. Leatherman. *Federal Income Taxation of Corporations and Partnerships.* Wolters Kluwer Law & Business in New York, 2014.

[12] S. G. Ficici and A. Bucci. Advanced tutorial on coevolution. In *GECCO (Companion)*, pages 3172–3204, 2007.

[13] GAO. Large partnerships: Growing population and complexity hinder effective irs audits. http://www.gao.gov/products/GAO-14-746T, 2014.

[14] GAO. Large partnerships: With growing number of partnerships, irs needs to improve audit efficiency. http://www.gao.gov/assets/670/665886.pdf, 2014.

[15] GAO. Partnerships and s corporations: Irs needs to improve information to address tax noncompliance. http://www.gao.gov/assets/670/663185.pdf, 2014.

[16] E. Hemberg, J. Rosen, G. Warner, S. Wijesinghe, , and U.-M. O'Reilly. Tax non-compliance detection using co-evolution of tax evasion risk and audit likelihood. In *The 15th International Conference on Artificial Intelligence and Law*, 2015.

[17] E. Hofmann, E. Hoelzl, and E. Kirchler. Preconditions of voluntary tax compliance: Knowledge and evaluation of taxation, norms, fairness, and motivation to cooperate. *Zeitschrift Fur Psychologie*, 216:209–217, 2008.

[18] IRS. Internal revenue manual, part 4 examining process, chapter 22 national research program, section 1 nrp overview, 2005. `http://www.irs.gov/irm/part4/irm_04-022-001.html`.

[19] C. Landstrom, S. Whatmore, and S. Lane. Learning through computer model improvisations. *Science, Technology & Human Values*, 38:678–700, 2013.

[20] R. M. Lipton. Bad dad – seventh circuit rejects distressed asset/debt shelter. *Journal of Taxation*, 119(5):236–240, 2013.

[21] D. McBarnet. Legtitimate rackets: Tax evasion, tax avoidance, and the boundaries of lity. *Critical Criminology: An International Journal*, 3(2):56–74, 1992.

[22] D. of the Treasury. The problem of corporate tax shelters: Discussion, analysis and legislative proposals. Technical report, Department of the Treasury, 1999.

[23] M. O'Neill and C. Ryan. *Grammatical evolution: evolutionary automatic programming in an arbitrary language*, volume 4. Springer, 2003.

[24] M. Pickhardt and A. Prinz. Behavioral dynamics of tax evasion–a survey. *Journal of Economic Psychology*, 40:1–19, 2014.

[25] J. Rosen, E. Hemberg, G. Warner, S. Wijesinghe, , and U.-M. O'Reilly. Computer aided tax evasion policy analysis: Directed search using autonomous agents. In *International Conference on Autonomous Agents and Multiagent Systems*, 2015.

[26] T. Rostain and M. C. Regan. *Confidence Games: Lawyers, Accountants and the Tax Shelter Industry*. Massachusetts Institute of Technology Press, 2014.

[27] P. Sancilio. Clarifying (or is it codifying) the "notably abstruse": Step transactions, economic substance, and the tax code. *Colombia Law Review*, pages 138–179, 2013.

[28] S. Shay. Interview with stephen shay, march 5, 2015. Personal Communication, March 2015.

[29] J. B. Sowell. The partnership anti-abuse rules: Where have we been and where are we going? Technical report, KPMG, 2011.

[30] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *J. Artif. Intell. Res.(JAIR)*, 21:63–100, 2004.

[31] F. Stonedahl, M. Wilkerson-Jerde, and U. Wilensky. Re-conveiving introductory computer science curricula through agent-based modeling. In *Proceedings of the EduMAS workshop at AAMAS 2009, May 12*, pages 63–70. AAMAS, 2009.

[32] W. M. Thomas. American jobs creation act of 2004. Became Public Law No: 108-357, 2004. H.R.4520.

[33] G. Warner, S. Wijesinghe, U. Marques, O. Badar, J. Rosen, E. Hemberg, and U.-M. O'Reilly. Modeling tax evasion with genetic algorithms. *The Economics of Governance*, 15(5), 2014.

[34] R. P. Wiegand and M. A. Potter. Robustness in cooperative coevolution. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 369–376. ACM, 2006.

[35] D. Wright Jr. Financial alchemy: How tax shelter promoters use financial products to bedevil the irs (and how the irs helps them), 2013.

[36] G. K. Yin. Getting serious about corporate tax shelters: Taking a lesson from history. *SMU Law Review*, 54(2), 2001.