Estimating Marginal Revenue/Cost Curves for Correspondence Audits

Ronald H. Hodge II, Alan H. Plumley, Kyle Richison, and Getaneh Yismaw,
Research, Analysis, and Statistics, Internal Revenue Service, *and*
Nicole Misek, Matt Olson, and H. Sanith Wijesinghe, MITRE Corporation

Tax agencies have long desired to allocate their resources so as to maximize the revenue they collect net of administrative costs.[1] It has been clear that net revenue is maximized when the marginal revenue/cost ratio is equalized across all potential activities; otherwise, net revenue could be increased by shifting resources from activities having low marginal revenue/cost to those that exhibit higher marginal revenue/cost. However, marginal revenue and marginal cost are usually not observed; they must be estimated, and that is often a challenge. As a result, many tax administrators rely on readily observable average revenue/cost metrics, which lead to sub-optimal outcomes.[2]

This paper provides empirical estimates of marginal revenue/cost functions for several important categories of correspondence audits of tax returns conducted by the Internal Revenue Service (IRS) for Tax Years 2006–2010, and explains the methodology we developed for making those estimates. We then use these curves to identify the allocation of resources among these audit categories for each of these historical years that would have maximized net direct revenue—suggesting that the same resources could have yielded an additional $190 million of direct enforcement revenue if they had been allocated differently during these years.

1. Introduction

It is quite easy and intuitive to derive the average revenue per case, the average cost per case, and the average revenue-to-cost ratio (often called "Return on Investment," or ROI) for a tax enforcement program. These metrics have the great advantage of being very straightforward to derive from observable data, and they have some use in managing enforcement programs and documenting results. However, none of these average measures provide the right basis for allocating scarce resources to the programs competing for those resources. That is, they cannot tell us *how much* of our resources we *should* devote to each program. In fact, devoting more of our budget to the programs exhibiting the highest average ROIs is often *not* the most cost-effective way to manage our resources. That is because our ultimate objective ought to be to maximize the net benefit we produce through our programs—that is, the total benefit minus the total cost—and the way to maximize net benefits is to equalize the *marginal* benefit/cost ratio across all programs; otherwise, net benefits could be increased by shifting resources from activities having low marginal benefit/cost to those that exhibit higher marginal benefit/cost.[3] One of the benefits produced by a tax agency is the revenue it collects—whether paid voluntarily or in response to enforcement actions.[4] This paper focuses solely on the revenue generated *directly*[5] by one enforcement program—correspondence audits.[6]

---

[1] Sometimes in addition to other objectives.

[2] See the Government Accountability Office report, "TAX GAP: IRS Could Significantly Increase Revenues by Better Targeting Enforcement Resources," GAO-13-151, December 2012.

[3] More accurately, the uniform marginal benefit/cost ratio would apply to any program that is not otherwise constrained. For example, it would presumably be possible to divert resources away from any program whose marginal benefit/cost ratio is currently *less* than the optimal uniform ratio. But if its ratio is *greater* than the optimal uniform ratio, and there are short-term or permanent constraints preventing the program from being expanded to the point at which its ratio is lowered to the uniform ratio, then once that constraint becomes binding, the marginal benefit/cost ratio would effectively drop to zero. At that point, it would be counter-productive (at least while the constraint remains binding) to add resources to that program, though it would be helpful to make progress on relieving the constraint, if possible, cost-effectively.

[4] We do not include any penalties or interest among the benefits since our objective is not to maximize them.

[5] If we had estimates of the associated changes in voluntary compliance that are induced indirectly by that program throughout the entire population, those estimates could be added to the direct revenue estimates to represent the full benefit of the program. Although that is a long-term goal of our research, it is reasonable in the meantime to assume that these indirect effects are the same for each type of correspondence audit, so that they wouldn't influence the optimal allocation across correspondence audit categories.

[6] Actually, we focus on just some of the major discretionary categories of correspondence audits.

The marginal revenue generated by a tax enforcement program generally declines as the level of effort (i.e., the budget) expended in that program increases.[7] At a given budget, the marginal revenue is the change in revenue associated with changing that level of effort a little bit. Likewise, the marginal cost is the corresponding change in cost that produces that marginal revenue. Generally, marginal effects are not observed in isolation; they must be estimated. The research presented in this paper is (to our knowledge) the first to develop empirical estimates of marginal direct revenue/cost ratios for an IRS program, and it reflects the following key features:

- Analyzing the entire population of completed audits in given categories;
- Estimating for each audit the full administrative cost to conduct the audit, assess additional tax, and collect the tax due;
- Accounting for the amount of additional tax actually collected—not simply the amount proposed by the auditors or the amount formally assessed after appeal or litigation; and
- Ranking actual and potential audits in a given category according to their operational priority (risk).

The remainder of the paper proceeds as follows: Section 2 describes the correspondence audit program at the IRS; Section 3 outlines our estimation methodology; Section 4 presents our empirical results; Section 5 illustrates how these estimates could be used to improve the allocation of the budget within the correspondence audit program, surveying the volatility of the curves over time; and Section 6 concludes.

2. The Correspondence Audit Program

The IRS audits[8] a small percentage of the tax returns filed in any given year. Most of these audits are initiated because the IRS perceives that the taxpayer may have misreported his tax obligation. There are three types of audits of individual tax returns, depending on the perceived nature of the noncompliance. The most complex issues require a face-to-face audit at the taxpayer's residence or place of business. Somewhat simpler cases are handled face-to-face in an IRS office. The simplest audits—typically requiring just documentation or responses to simple inquiries about a limited number of issues—are conducted through correspondence. Since Fiscal Year 2006, just over three-quarters of all of the individual income tax audits that were completed were conducted through correspondence, and they accounted for nearly 60 percent of the tax adjustments proposed by all individual income tax audits each year (Table 1).[9] Correspondence audits tend to cost less—both in terms of IRS direct costs and also in terms of the burden placed on the taxpayers—so they are often quite cost-effective compared with face-to-face audits; hence the reliance on correspondence audits, which also underscores the importance of allocating our resources to the various categories of correspondence audits in the most cost-effective mix possible. Correspondence audits are categorized into "projects", with each project devoted to one or a small number of specific lines on the tax return (e.g., specific sources of income, specific deduction items, a tax credit, or some combination of these), or to the special issues that are often present on returns filed late. Each project, then, consists of a fairly homogeneous set of returns. The returns within a given project that are actually selected for correspondence audit are even more homogeneous because they all must meet a set of project-specific selection criteria that suggest the possibility of misreporting. This first set of selection criteria (which are generally binary, in that they have a yes/no quality) typically produces more workload than the program has the resources to pursue. So, the program selects from among the returns that meet these initial screening criteria by prioritizing them according to a non-binary (i.e., continuously varying) indicator or computed risk score.

At the conclusion of any type of audit, the examiner proposes the final tax adjustment (if any) to the taxpayer. If the taxpayer agrees with this "recommended" adjustment, that amount is formally assessed and becomes legally due. If the taxpayer disagrees with some or all of the recommended amount, the taxpayer must challenge it either through administrative appeal with the IRS or through litigation, both of which will result in either an assessment of some or all of the recommended amount or an assessment of zero. Regardless of how or when an assessment is made, it is legally due. The assessed amount is often paid in full immediately (after a formal notice is sent to the taxpayer) or in installments, but sometimes collecting the tax due requires the IRS to devote additional resources to follow up with the taxpayer, and even to apply stronger collection tools such as liens or levies.

[7] This is because most enforcement programs are successful in giving priority to cases that are at least somewhat more cost-effective to work than others. If the tax agency had no way of identifying in advance which cases would be more cost-effective than others, it would inevitably select cases randomly, whereupon the marginal revenue/cost ratio would be the same as the overall average—a constant.

[8] Technically, the "audits" should be referred to as *examinations*, since they are not comprehensive audits of the taxpayer's tax obligation.

[9] *Internal Revenue Service Data Book,* Table 9a; available each fiscal year on http://www.irs.gov/uac/SOI-Tax-Stats-IRS-Data-Book.

TABLE 1. Correspondence Examinations as a Share of all Examinations, FY 2006–2014

Fiscal Year	Returns Examined		Percent Correspondence	Recommended Additional Tax ($ Millions)		Percent Correspondence
	Field	Correspondence		Field	Correspondence	
2006	302,785	981,165	76.4%	5,433	7,612	58.4%
2007	311,339	1,073,224	77.5%	6,357	9,348	59.5%
2008	310,429	1,081,152	77.7%	5,945	6,518	52.3%
2009	326,249	1,099,639	77.1%	7,145	7,796	52.2%
2010	342,762	1,238,632	78.3%	6,899	8,168	54.2%
2011	391,621	1,173,069	75.0%	5,947	8,705	59.4%
2012	359,750	1,122,216	75.7%	6,062	9,249	60.4%
2013	344,152	1,060,779	75.5%	5,594	8,455	60.2%
2014	291,643	950,836	76.5%	5,026	6,859	57.7%

NOTE: Roughly 500,000 correspondence audits relate to the Earned Income Tax Credit.
Source: IRS Data Books, Fiscal Years 2006–2014.

3. Estimation Methodology

The ultimate objective of any tax enforcement program ought to be to maximize its net benefits, which we can simplify in our context to mean maximizing net direct revenue (the total revenue actually collected minus the total cost of the audit, appeals, litigation, and collection processes applied to collect that revenue).[10] It's not cost-effective to maximize net recommendations or net assessments; if the costs incurred on an audit produce a recommended or assessed tax change, but they don't actually produce net revenue collected, then the costs expended on that case were wasted.[11] Therefore, the methodology presented in this paper estimates marginal *revenue* (the amount of tax ultimately collected or paid) and marginal *ultimate* cost (the cost of the full audit, appeals, litigation, and collection life cycle of the cases).

Our methodology followed nine basic steps:

1. Segment the Population: We divide all completed audits for a given year into groups that are somewhat homogeneous with respect to taxpayer and agency behavior. For correspondence audits, the projects were the appropriate groups since each project reflects a unique compliance behavior and has its own criteria for selecting tax returns for audit. For some projects, we further segment the returns into High Income and Low Income sub-groups.

2. Identify Revenue Collected: For each audit, we identify the amount of tax[12] that was eventually collected. Since this ultimate disposition of the case can take some time, it is best to focus on a year for which the collection process has run its course. Fortunately, few correspondence audits take an extended amount of time to fully close.

3. Estimate Cost: For each audit, we also estimate the total cost of the entire examination, appeals, litigation, notice, and collection process that was incurred to collect the revenue for that case. IRS enforcement data include the number of hours (by type of employee) spent in the examination, appeals, and litigation steps (if applicable),

[10] More generally, the benefits should include such things as the potential increase in revenue paid voluntarily in the general population due to the *indirect* (possibly deterrent) effect of the enforcement, and the costs should include any unnecessary monetary and non-monetary compliance costs borne by taxpayers and third parties in connection with the enforcement (routine, necessary post-filing compliance costs borne by taxpayers may serve as a deterrent, prompting taxpayers to avoid enforcement, which is a positive outcome that shouldn't be treated as a cost that nets against benefits). However, those non-direct benefits and costs are very difficult to quantify, and it is not clear how the various components of benefits and costs should be weighted relative to each other. Nonetheless, direct revenue and costs are the foundational components, and a tax enforcement program could maximize net direct revenue subject to reasonable constraints intended to account for the missing components from the ideal objective function.

[11] Some people might argue that conducting such an audit may well have an indirect deterrent value, promoting better voluntary compliance in the general population, but that would likely be more true if the case produced net revenue. And even if two cases produced the same indirect effect, one that produced no net direct revenue would clearly be less advantageous than one that did produce net direct revenue. Others may contend that not auditing returns we expect are understating their tax, but that we anticipate will not produce any net revenue, might prompt some taxpayers to display characteristics that could plausibly be expected to make them appear to be unlikely to pay a tax assessment without a lot of additional effort on the part of the IRS. It's possible that assessing less than the full recommended amount in the appeals or litigation process, or unsuccessful collection efforts might deter such taxpayers (or others in the population) from driving up the costs of IRS enforcement programs or reducing the collectability of enforcement assessments, but it seems more probable that the reverse is true as taxpayers may perceive that they can end up paying much less than the amount originally recommended by appealing the adjustment and dragging out the collection process.

[12] We excluded any interest and penalties that were paid, since we ought not let these become an incentive for us to delay the collection of tax. We also excluded non-enforcement revenue where it existed.

so it was a straightforward matter to estimate the cost of these steps using known hourly costs for each type of enforcement employee. Likewise, we have data indicating the number and type of automated notices sent to the taxpayer demanding payment of the final amount that was assessed, so it was easy to estimate that component of the cost. However, since the hours spent in the collection process are not captured separately for each tax return, we had to estimate that component of the overall cost for each audit. Fortunately, very few correspondence audits require such collection efforts. So for those that did, we applied the overall average cost per case of the collection process, accounting for the Automated Collection System and the Field Collection functions separately, depending on which of these a given audit case required.

4. Sort the Observations According to Priority: All of the returns in the database of completed audits met the basic project-specific selection criteria, but they differed with respect to a non-binary (i.e., continuously varying) indicator, as discussed in Section 2. Returns with a high value for the indicator were selected for audit, but the same indicator implicitly assigns a priority to each return, since those with the lowest value would not have been selected for audit had resources been tighter. So, we sorted the completed audits in declining order of the relevant indicator.

5. Compute Cumulative Revenue and Cumulative Cost: For each tax return in a given segment (i.e., project) in the given year, we computed the total revenue that would have been collected had that audit been the lowest priority audit to have been conducted. The revenue collected from each successive audit in the sorted file was added to the cumulative revenue up through the previous audit in the sorted file. The same procedure was used to compute the cumulative cost for each audit in the sorted file.

6. Plot Cumulative Revenue vs. Cumulative Cost: For each segment for a given year, we construct a plot of cumulative revenue vs. cumulative cost. This typically exhibits a slightly curved pattern, with the higher priority audits (on the left of the graph—at low levels of cumulative cost) having a higher revenue/cost than those with lower priority (on the right of the graph).

7. Fit a Curve: We used regression analysis to fit a curve through the observations in the plot constructed above. We found that a simple quadratic specification fit very well for most of the projects (which has simplifying benefits for the next step). In this case, we chose the following functional form:

$$\text{Cumulative Revenue} = a(\text{Cumulative Cost}) + b(\text{Cumulative Cost})^2 + \varepsilon \quad (1)$$

where a and b are the parameters to be estimated and ε is a disturbance term. No constant (intercept) term was included since cumulative revenue should be zero when cumulative cost is zero.

For some projects, a Power curve provided a better fit to the data. This has the following functional form:

$$\text{Cumulative Revenue} = e^{a}\,(\text{Cumulative Cost})^{b} + \varepsilon_0 \quad \text{or}$$

$$\ln(\text{Cumulative Revenue}) = a + b\ln(\text{Cumulative Cost}) + \varepsilon \quad (2)$$

where a and b are again the parameters to be estimated and ε is a disturbance term.

8. Derive Marginal Revenue/Cost: We can easily derive Marginal Revenue/Cost (MR/MC) as a function of Cumulative Cost (i.e., budget) by taking the first derivative of equations (1) and (2) with respect to Cumulative Cost. The results are:

$$\text{Quadratic: MR/MC} = a + 2b(\text{Cumulative Cost}) \quad (3)$$

$$\text{Power: MR/MC} = e^{a}\, b(\text{Cumulative Cost})^{b-1} \quad (4)$$

Since we specified equation (1) as a quadratic, equation (3) is simply a straight line, with intercept a and a slope of $2b$. Hence, a should be positive, as it represents the revenue/cost of the first (highest priority) audit. And since we expect declining MR/MC, b should be negative. The value of equation (3) at a given level of Cumulative Cost is the slope of the fitted curve (equation (1)) at that level of resources. Given the nature of Equation (2), Equation (4) is nonlinear, approaching the MR/MC-axis asymptotically, and approaching a horizontal line (relatively constant MR/MC) at high levels of cost. That is, for declining MR/MC, a should be positive, and b should be between 0 and 1.

9. Derive the Optimal Allocation:[13] As indicated earlier, the allocation that maximizes net direct revenue is the one in which the marginal revenue/cost ratio is the same across all segments competing against each other for resources. A straightforward way to do this is to transform equations (3) and (4) to express Cumulative Cost (i.e., budget) as a function of the marginal revenue/cost (MR/MC) ratio, as follows:

$$\text{Quadratic: Budget} = (\text{MR/MC} - a) / 2b \tag{5}$$

$$\text{Power: Budget} = \left[\frac{\text{MR/MC}}{be^{a}}\right]^{1/b-1} \tag{6}$$

We can then use equations (5) and (6) to calculate the budget required for each segment as a function of a common MR/MC ratio (using the estimated parameters *a* and *b* relevant to each segment), and therefore the total budget across all segments as a function of the common MR/MC ratio. The optimal allocation is the one in which the total budget estimated in this way is the same as the budget actually available (i.e., the overall cumulative cost).[14]

4. Empirical Results

We applied this methodology to the following seven correspondence audit project categories:

TABLE 2. Correspondence Audit Categories Studied

Project	Issue(s) Addressed
C1	A sole proprietor issue reported on Schedule C
A1-Lo	Various items claimed on Schedule A (Lo means low income taxpayers) (Hi means high income taxpayers)
A1-Hi	
A2	
A3	
A4	
O	Another Form 1040 issue

These projects account for many—but not all—of the correspondence audits conducted during these years as an illustration of the methodology. Table 3 provides an overview of the data for these seven projects for Tax Years 2006–2010, the most recent years available for analysis. Notice that the number of audits conducted and the revenue and cost varied significantly across these projects. Project O had the highest volumes and the highest average revenue/cost ratio, but it's not clear from these average metrics how to allocate the budget across these seven project categories. That will depend on the *marginal* revenue/cost functions, which we estimated using the methodology described in Section 3. Figures 1 through 7 show the basic plots of cumulative revenue vs. cumulative cost using the raw data, as well as the fitted curves through those data and the overall average revenue/cost lines for Tax Year 2006. Notice that they all display the expected curvature (diminishing marginal revenue/cost), although it is typically rather modest (presumably due to the rather homogeneous nature of the projects and the simple method used to assign priority to the cases).

[13] We use the term "optimal" in the narrow context in which our simplified objective is to maximize net direct revenue (i.e., the revenue collected directly from the audits minus the full administrative cost to identify, assess, and collect that revenue). This may not be optimal in the ultimate context in which we account for all other benefits and costs to the IRS, taxpayers, and third parties, but this is a necessary starting point. Moreover, it is possible to impose constraints in this simple framework that limit the extent to which segments (i.e., projects) can be expanded or contracted. Some of these constraints may be known and quantifiable workload or resource limitations, while others may be subjective rules (such as minimum coverage constraints) that attempt to account for the likelihood and impact of indirect benefits.

[14] This procedure would need to be modified to account for any constraints imposed on the expansion or contraction of the audit program in any of the segments. In the absence of such a constraint, a segment whose marginal revenue/cost ratio is below the optimal ratio at all budget levels would not have any audits conducted at all, and a segment whose marginal revenue/cost ratio is above the optimal ratio even for the last potential audit in the population would be given the budget to conduct audits on all possible returns.

TABLE 3. Summary Data for Selected Correspondence Audit Projects, Tax Years 2006–2010

Project	Tax Period	Number of Exams Completed	Total Revenue Collected ($M)	Total Administrative Cost ($M)	Average Revenue/ Cost	Average Revenue/ Exam	Average Cost/Exam
C1	2006	10,706	$28.2	$4.0	7.0	$2,636	$374
	2007	24,740	$68.8	$8.7	7.9	$2,781	$351
	2008	19,114	$48.2	$5.3	9.1	$2,522	$277
	2009	23,185	$61.6	$7.4	8.3	$2,657	$321
	2010	29,495	$40.1	$5.9	6.8	$1,358	$200
A1 Low Income	2006	67,621	$159.4	$13.4	11.9	$2,357	$198
	2007	58,615	$126.6	$12.0	10.6	$2,161	$204
	2008	27,803	$64.4	$5.2	12.3	$2,317	$188
	2009	41,528	$72.0	$8.1	8.9	$1,734	$196
	2010	31,614	$36.9	$5.7	6.5	$1,169	$181
A1 High Income	2006	9,895	$19.3	$2.4	8.1	$1,947	$240
	2007	9,459	$27.1	$2.7	10.0	$2,870	$286
	2008	7,813	$23.9	$2.5	9.5	$3,063	$323
	2009	3,428	$2.1	$0.7	3.2	$623	$196
	2010	3,080	$3.9	$0.6	6.0	$1,261	$210
A2	2006	30,541	$67.7	$7.1	9.5	$2,216	$233
	2007	27,639	$62.2	$6.8	9.2	$2,252	$245
	2008	2,480	$5.1	$0.4	14.2	$2,070	$146
	2009	8,838	$15.1	$2.0	7.6	$1,714	$226
	2010	2,135	$3.3	$0.5	6.6	$1,540	$233
A3	2006	21,298	$22.3	$4.3	5.2	$1,048	$201
	2007	10,961	$9.2	$2.0	4.7	$840	$179
	2008	21,011	$16.2	$3.4	4.7	$771	$163
	2009	13,043	$8.3	$1.6	5.3	$634	$119
	2010	12,768	$7.3	$1.8	4.1	$570	$139
A4	2006	13,919	$37.5	$3.0	12.6	$2,696	$214
	2007	41,069	$86.8	$9.1	9.6	$2,114	$220
	2008	23,885	$63.1	$5.6	11.2	$2,641	$235
	2009	17,942	$48.9	$4.9	10.1	$2,723	$271
	2010	31,799	$50.7	$5.9	8.6	$1,595	$184
O	2006	202,575	$287.1	$37.0	7.8	$1,417	$183
	2007	116,555	$177.4	$21.2	8.4	$1,522	$182
	2008	105,503	$96.3	$18.0	5.4	$913	$170
	2009	78,994	$79.2	$12.3	6.4	$1,003	$156
	2010	29,711	$15.6	$3.7	4.2	$525	$126

FIGURE 1. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code C1, Tax Year 2006

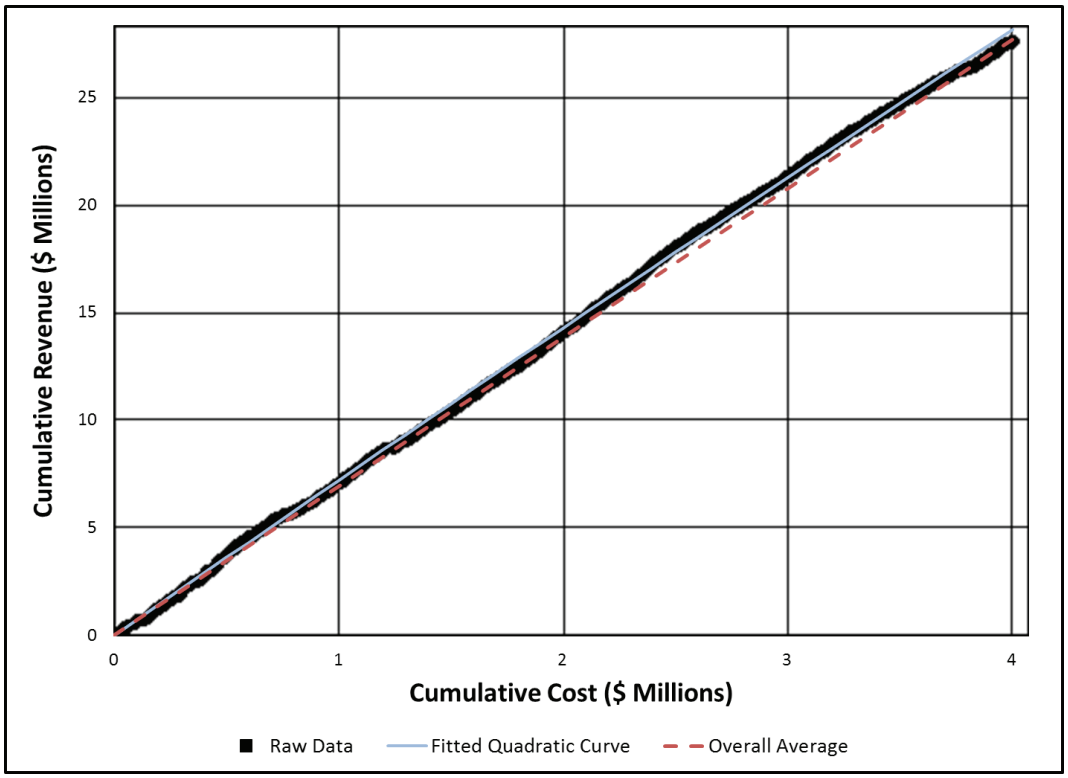

FIGURE 2. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code A1-Low Income, Tax Year 2006

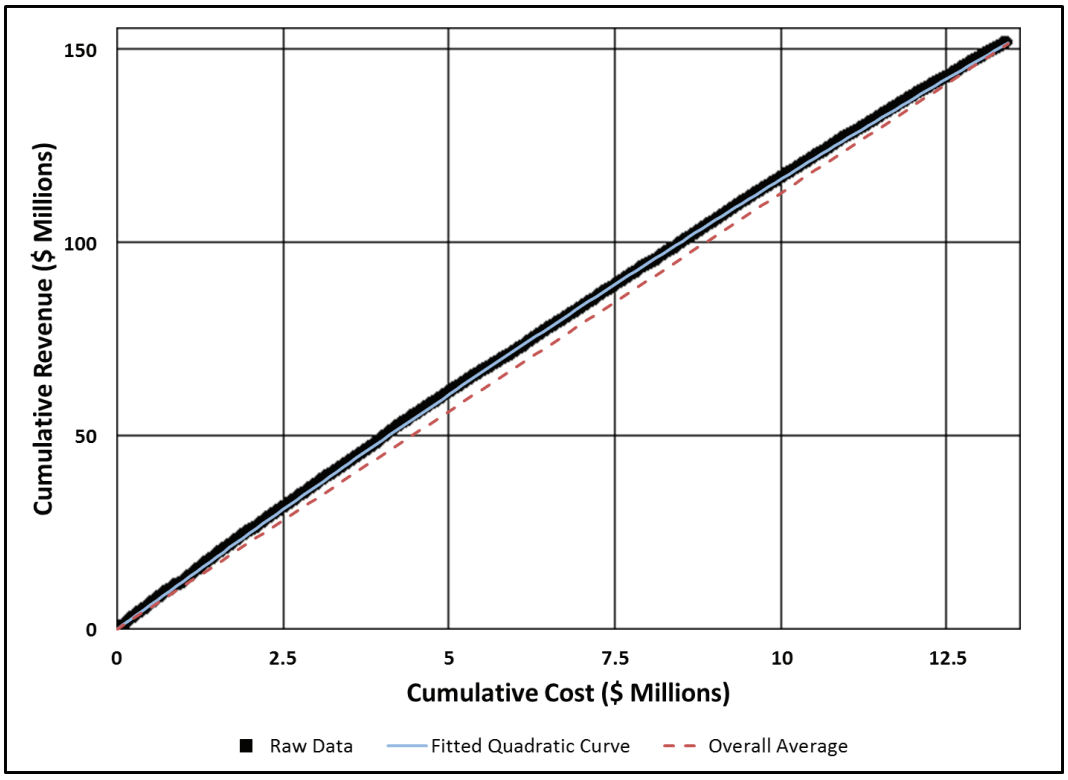

FIGURE 3. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code A1-High Income, Tax Year 2006

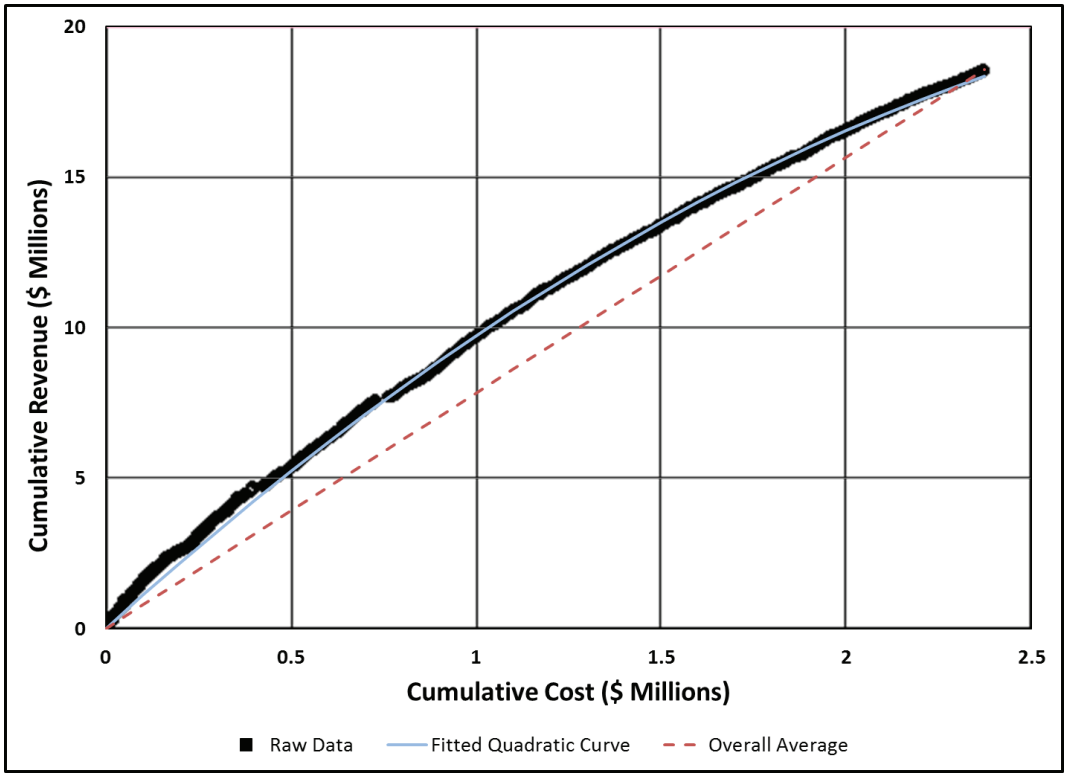

FIGURE 4. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code A2, Tax Year 2006

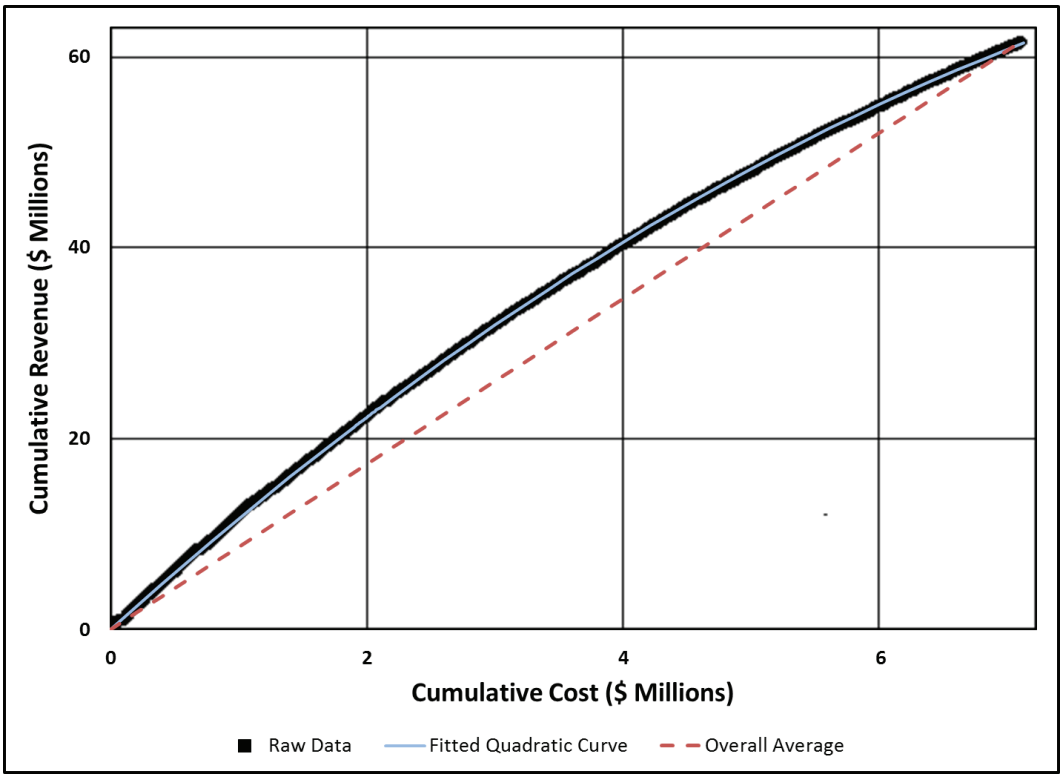

FIGURE 5. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code A3, Tax Year 2006

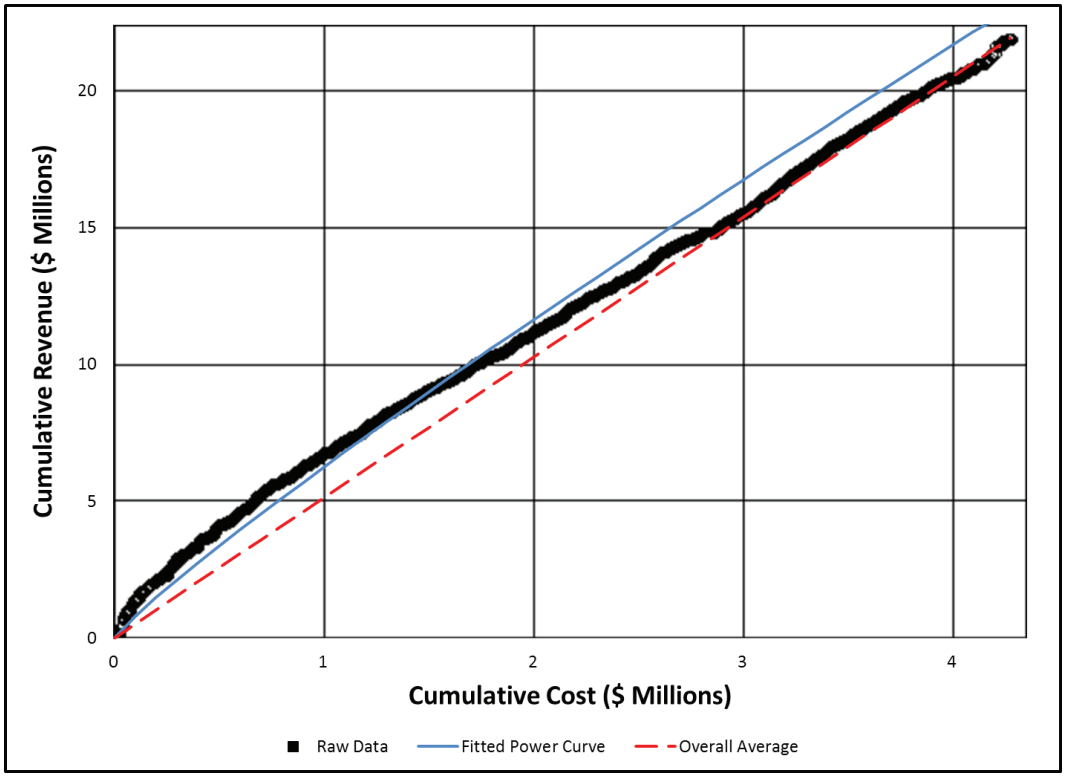

FIGURE 6. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code A4, Tax Year 2006

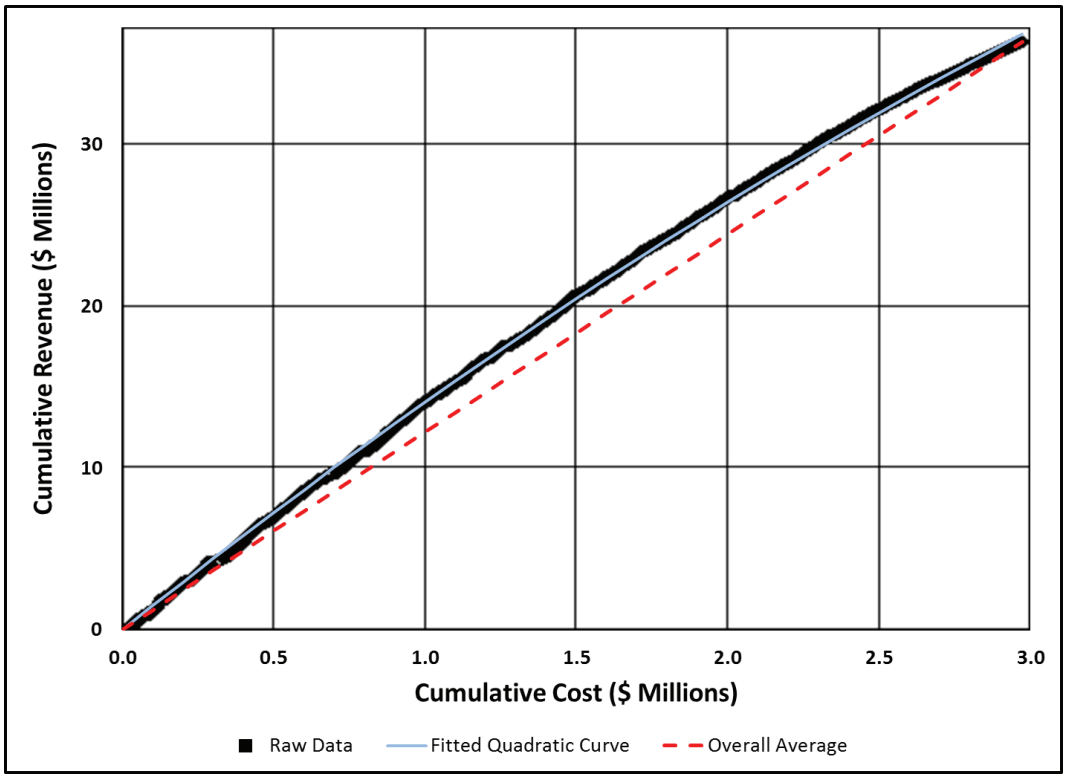

FIGURE 7. Plot of Sorted Data and Average and Fitted Curves of Revenue vs. Cost, Correspondence Audit Project Code O, Tax Year 2006

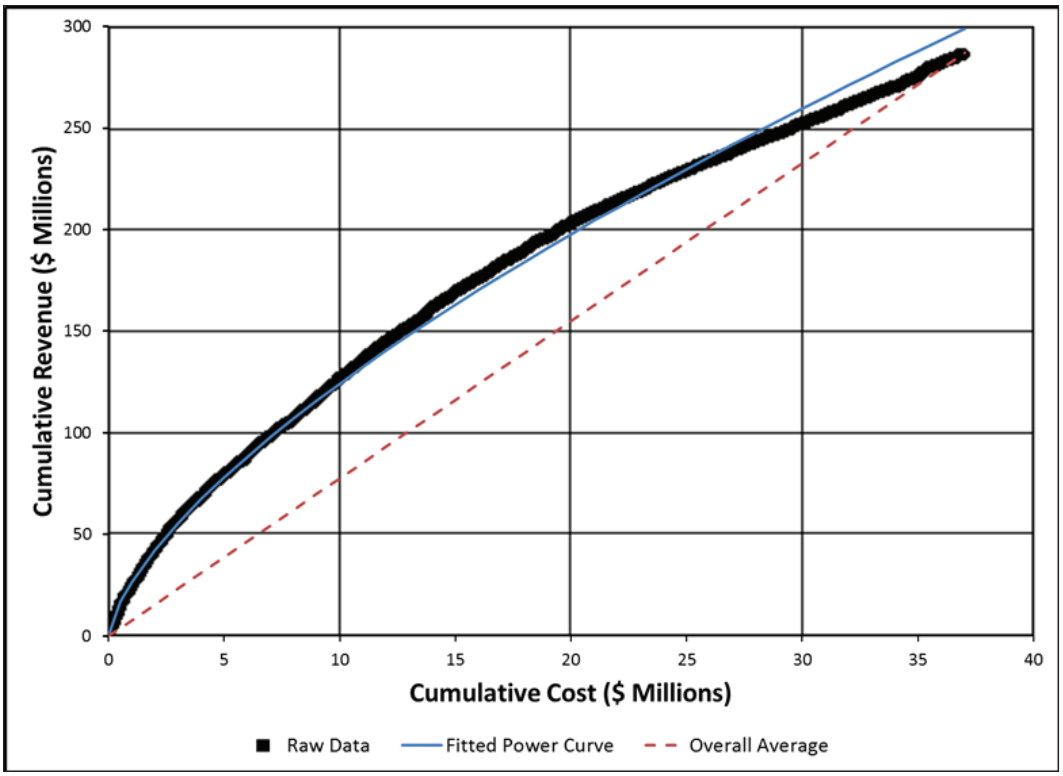

Notice that some of the curves (particularly Project Codes A2 and O) demonstrate strongly declining marginal revenue/cost (as reflected in their curvature), while others (e.g., Project Codes C1 and A1-Low) differ very little from the overall average line. In all of the projects, however, it may be possible to identify an improved prioritization metric, which would allow us to select much more cost-effective returns to audit than at present, resulting in more strongly declining revenue/cost and much higher overall revenue.[15]

The regression results for the fitted curves for Tax Year 2006 are presented in Table 4. All of the parameters are highly significant (due to the large number of observations), and have the expected signs. The results for the other years are similar, except that the parameter *b* for Project Code A1-Low was slightly positive in Tax Year 2008 (indicating a slightly *increasing* marginal revenue/cost due to the weakness of the prioritization variable).

5. Resource Allocation Implications

Applying equations (3) and (4), we can easily derive the marginal revenue/cost for each project as a function of the cumulative cost (or budget) allocated to it. This is merely the slope of the fitted curve at each point. Figure 8 plots all seven marginal curves together, and Table 5 provides the numerical results. Given a total budget of $71.2 million, the maximum net direct revenue that could have been generated by these seven projects was $597.4 million or $63.5 million more than was actually generated. This outcome would have been produced if the budget had been allocated to these projects such that they all shared a marginal revenue/cost of 6.51. This allocation would have shifted resources from Projects O, A1-Hi, A2, and A3 to Projects C1, A1-Lo, and A4. For the most part, the employees who audit returns in these projects are interchangeable, and there appear to be ample returns that could have been audited in Project A1-Lo, so this allocation appears to have been feasible.[16] It is not certain that had Project A1-Lo been expanded to this extent the estimated marginal revenue/cost function would have continued along the same straight line, but the strong relationship between revenue and cost suggests that this assumption is a strong basis for resource allocation.

[15] This will be the subject of future research.

[16] If this were not true in the short term, steps could be taken over time to move toward the optimal allocation, to the extent this is cost-effective.

TABLE 4. Regression Results for the Estimated Curves, Tax Year 2006
(Dependent Variable = Cumulative Revenue, CR)

Project	Functional Form (where C = cumulative cost)	a	b	Adjusted R-squared
C1	$CR = aC + bC^2$	7.23029 (2021.77)	-5.21147E-8 (-45.91)	0.9999
A1-Lo	$CR = aC + bC^2$	12.70311 (16684.8)	-1.01473E-7 (-1399.8)	1.0000
A1-Hi	$CR = aC + bC^2$	11.17570 (2228.35)	-1.45E-6 (-534.09)	0.9998
A2	$CR = aC + bC^2$	12.09733 (13223.5)	-4.87995E-7 (-2974.0)	1.0000
A3	$CR = e^aC^b$	3.23558 (117.31)	0.89817 (465.98)	0.9107
A4	$CR = aC + bC^2$	14.91579 (3639.92)	-8.55108E-7 (-495.50)	0.9999
O	$CR = e^aC^b$	7.83004 (2623.44)	0.67055 (3740.45)	0.9857

NOTE: t-statistics in parentheses

FIGURE 8. Estimated Marginal Revenue/Cost Curves, Selected Projects, Tax Year 2006

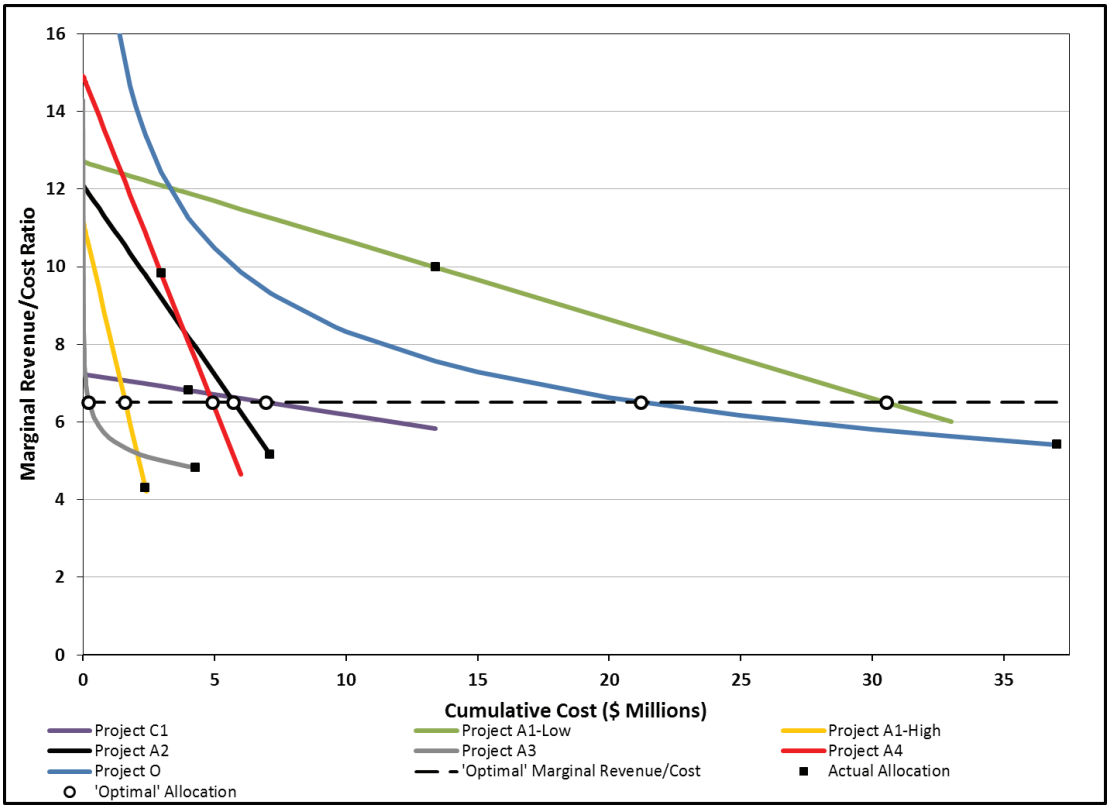

TABLE 5. Actual vs. Optimal* Revenue and Cost, Tax Year 2006 ($ in Millions)

Project	Actual Allocation					Optimal* Allocation					Change in Cost	Change in Revenue
	Total Revenue	Total Cost	Net Revenue	Average Rev/Cost	Marginal Rev/Cost	Total Revenue	Total Cost	Net Revenue	Average Rev/Cost	Marginal Rev/Cost		
C1	$27.6	$4.0	$23.6	6.90	6.81	$47.7	$7.0	$40.8	$6.87	6.51	$2.9	$20.1
A1-Lo	$151.9	$13.4	$138.5	11.33	9.98	$293.3	$30.5	$262.7	$9.60	6.51	$17.1	$141.4
A1-Hi	$18.6	$2.4	$16.2	7.83	4.29	$14.2	$1.6	$12.6	$8.84	6.51	-$0.8	-$4.3
A2	$61.7	$7.1	$54.6	8.67	5.16	$53.3	$5.7	$47.6	$9.30	6.51	-$1.4	-$8.4
A3	$21.9	$4.3	$17.6	5.12	4.82	$1.6	$0.2	$1.4	$7.24	6.51	-$4.1	-$20.3
A4	$36.3	$3.0	$33.4	12.22	9.83	$52.7	$4.9	$47.8	$10.71	6.51	$1.9	$16.3
O	$287.1	$37.0	$250.0	7.75	5.41	$205.7	$21.2	$184.5	$9.70	6.51	-$15.8	-$81.4
Total	**$605.1**	**$71.2**	**$533.9**	**8.50**	**N/A**	**$668.6**	**$71.2**	**$597.4**	**$9.39**	**6.51**	**$0.0**	**$63.5**

*Optimal only in the sense of maximizing direct revenue in an unconstrained setting, without accounting for indirect effects

Table 5 also shows that, at least for these projects, the average revenue/cost ratios under the actual allocation did not have the same ranking as the corresponding marginal revenue/cost ratios, and the optimal allocation was not at all proportional to or consistent with those average revenue/cost ratios. Most importantly, if the same resources had been allocated according to the *average* ratios, the largest share of the budget would presumably have been devoted to Project A4 (since its ratio of 12.22 was highest), but the optimal allocation (unlike the actual allocation) devotes the largest share of the budget to Project A1-Lo, which also had the highest marginal revenue/cost ratio in the actual allocation. The key is that without the marginal revenue/cost framework, even though decision-makers might want to shift resources from Projects C1 and A3 (given their low average ratios), they would not know *how much* of each budget to redirect to Projects A4 or A1-Lo; indeed they likely would have taken resources *from* Project C1 instead of *adding* resources to it, and they likely would have favored Project A4 over Project A1-Lo (given their relative average ratios).

Year-to-Year Fluctuations

Ideally, the relationship between revenue and cost is relatively stable from year to year. However, this seems not to have been true in general. For example, Figure 9 shows how the cumulative revenue vs. cumulative cost plot for Project A4 varies across the five years of this study. While there was some consistency, there was also significant variation.

FIGURE 9. Cumulative Revenue vs. Cumulative Cost Curves, Project A4, Tax Years 2006–2010

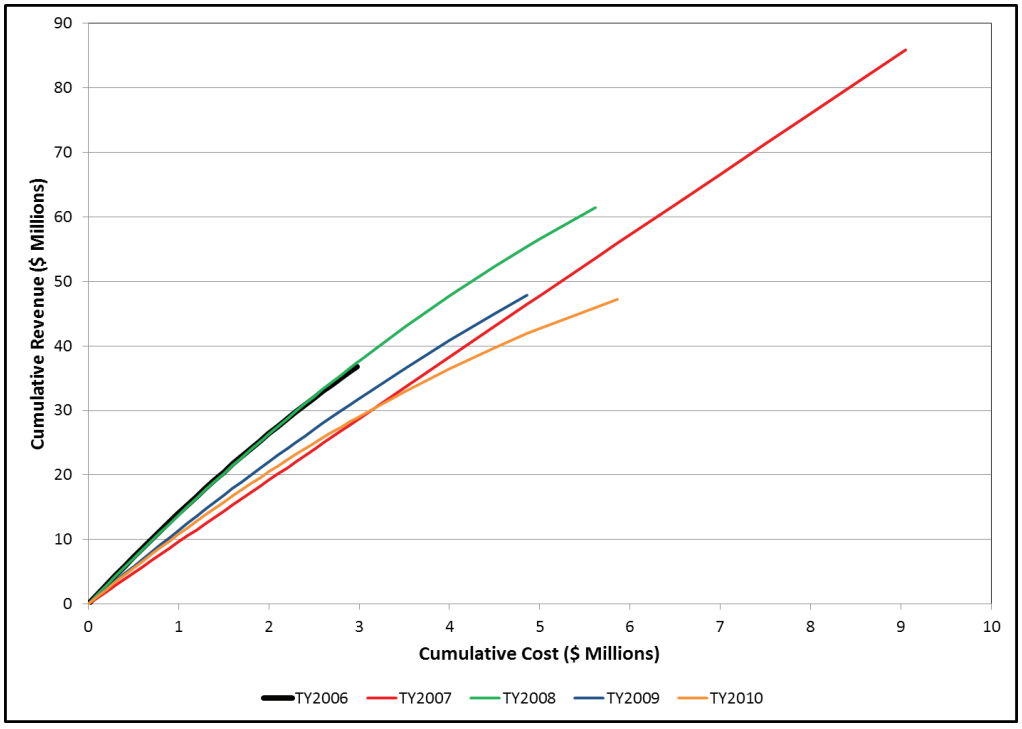

There was even more variation in Project A1-Hi, as shown in Figure 10. In both cases, the budget (total cost) varied widely across these years, but the curves themselves varied, too.[17] One might expect that the relationship between cumulative revenue and cumulative cost would remain fairly stable, and that a change in the budget would represent moving to a different location on the curve. But this appears not to be the case, making it challenging to use historical audit results to guide resource allocation in the current year.

FIGURE 10. Cumulative Revenue vs. Cumulative Cost Curves, Project A1-Hi, Tax Years 2006–2010

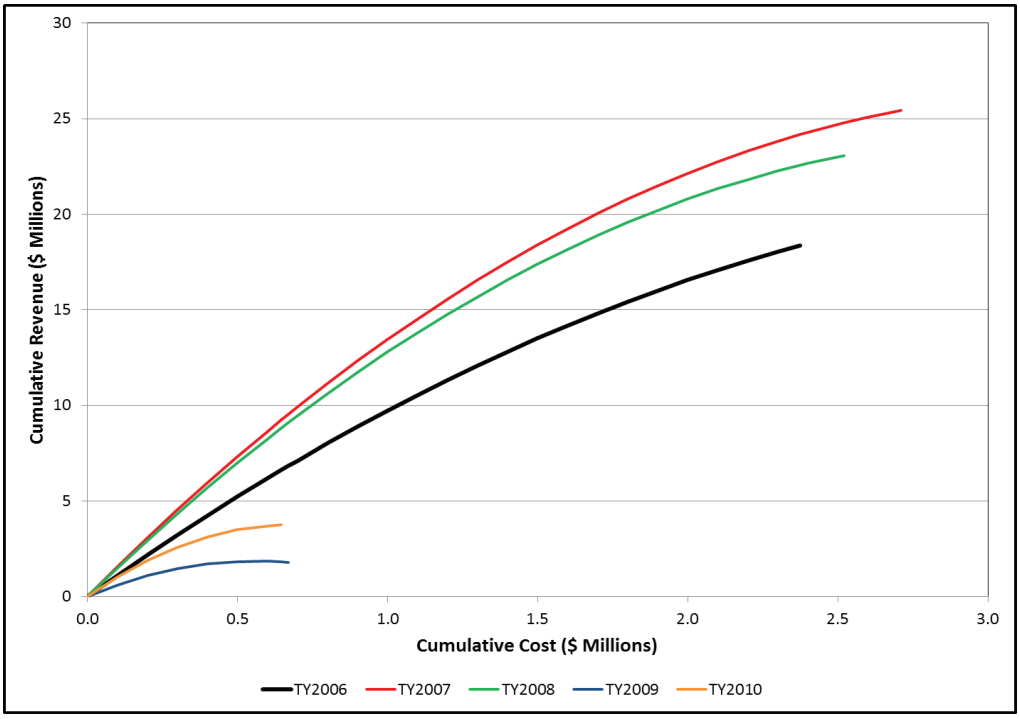

The variation over time is also seen in Table 6, which presents the change in revenue and the change in cost (budget allocation) for each project and year had the budget been allocated "optimally" with perfect knowledge of the marginal revenue/cost curves for the year in question. Obviously, the changing curves produce a very different "optimal" mix of projects each year. Further evidence of the time-sensitive nature of these relationships is given in Table 7, which corresponds to the estimates summarized in Table 6. Notice that the overall budget for these seven projects declined steadily from $71.2 million to $25.1 million, and the resulting revenue declined from $605.1 million to $141.9 million. If the budget had been allocated across these projects so as to maximize direct revenue, taking into account the nonlinear relationship between revenue and cost, one would expect that both the average and marginal revenue/cost ratios would increase as the budget decreased, but this did not happen.

There are undoubtedly many reasons for these observations, including the following (not necessarily listed in the order of their impact):

- In reality, the IRS does not allocate its resources solely to maximize direct revenue. Other objectives—such as maximizing the dollars assessed (rather than collected), minimizing the no-change rate, minimizing the time an audit is open, stabilizing mail volumes, and maximizing the presumed impact of the audits on the voluntary compliance of the general population—are routinely pursued, as well. Hence, the historical data upon which our curves are based are not always consistent with the revenue-maximizing assumptions we make. There are other factors that influence return selection for correspondence audits. For example, the First-Time Homebuyer Credit took effect during this period, causing correspondence audit resources to be diverted from regular projects for a couple of years to monitor this new compliance challenge.

[17] Similar variations were exhibited by the other projects in this study over this time period.

- Taxpayers' compliance behavior undoubtedly changes over time. This may be particularly true on any specific tax return line item. Since correspondence audits focus on a small number of related return line items (often just one line item), the results of these audits tend to be very sensitive to trends in taxpayer behavior.

TABLE 6. Change in Revenue and Cost: "Optimal" Allocation vs. Actual Allocation Among Selected Correspondence Audit Projects, Tax Years 2006–2010 ($ Millions)

Project	TY2006	TY2007	TY2008	TY2009	TY2010	Total
Change in Allocation of Budget						
C1	$2.9	-$5.3	$3.3	$2.0	-$0.6	$2.3
A1-Low	$17.1	$0.1	$0.5	$5.3	-$0.8	$22.3
A1-High	-$0.8	-$1.2	-$0.1	-$0.6	-$0.2	-$2.8
A2	-$1.4	-$3.6	$0.1	$0.0	-$0.1	-$5.0
A3	-$4.1	-$1.9	-$1.5	-$1.3	$2.8	-$5.9
A4	$1.9	$26.2	$3.0	$3.0	$0.6	$34.7
O	-$15.8	-$14.3	-$5.3	-$8.5	-$1.6	-$45.5
Total	**$0.0**	**$0.0**	**$0.0**	**$0.0**	**$0.0**	**$0.0**
Change in Revenue						
C1	$20.1	-$33.3	$16.5	$13.0	-$2.8	$13.5
A1-Low	$141.4	$2.3	$5.8	$32.7	$0.4	$182.6
A1-High	-$4.3	-$7.2	-$0.9	-$1.4	-$0.5	-$14.3
A2	-$8.4	-$24.8	$0.3	$0.3	-$0.5	-$33.1
A3	-$20.3	-$8.3	-$6.4	-$5.1	$12.0	-$28.0
A4	$16.3	$234.3	$17.0	$20.2	$2.3	$290.2
O	-$81.4	-$90.1	-$12.6	-$33.7	-$3.8	-$221.5
Total	**$63.5**	**$72.8**	**$19.9**	**$26.1**	**$7.1**	**$189.4**

TABLE 7. Change in Average and "Optimal" Revenue/Cost Ratios Over Time, Selected Correspondence Audit Projects, Tax Years 2006–2010

	TY2006	TY2007	TY2008	TY2009	TY2010
Total Cost ($ Millions)	$71.2	$62.3	$40.4	$37.0	$25.1
Total Original Revenue ($ Millions)	$605.1	$543.2	$307.9	$273.2	$141.9
Overall Original **Average** R/C	8.50	8.71	7.61	7.39	5.65
"Optimal" MR/MC ratio	6.51	8.43	3.69	5.39	4.03

6. Conclusion

This simple exercise involving just seven projects within the correspondence audit function illustrates the great potential for using this framework to increase net direct revenue through a reallocation of resources. Although the optimal allocation would be much more complicated to determine once we have estimated similar marginal revenue/cost functions for all other correspondence audit projects—and particularly for other enforcement programs—the basic approach would be the same.

In addition to paving the way for estimating similar marginal curves for other enforcement activities, this research is already leading to improved risk assessment formulas for prioritizing returns to audit. In the context of shrinking budgets and demands for improved revenue collection, this research promises to help the IRS move closer to the optimal allocation of its resources. However, the volatility of the estimated curves over time makes it very challenging to have confidence that using historical audit data to allocate resources in the current year will guide us close to the "optimal" allocation. Further research is under way to tackle this challenge.