Optimizing Policy Targeting with Machine Learning: Evidence from Pakistani Audits

Nicholas Lacoste[1] and Zehra Farooq[2]

[1]Tulane University
[2]Federal Board of Revenue, Pakistan

December 4, 2024

Abstract

We provide a framework for optimizing the distribution of policy treatment by combining machine learning techniques for the estimation of individualized causal responses with sufficient statistics for relative welfare implications of alternative distributions. This framework is applied to the setting of audit policy optimization in Pakistan. Specifically, we define a model which derives the Marginal Value of Public Funds (MVPF) in terms of three estimable causal effects of individuals in response to an audit: the net-present value of long-run tax liabilities, taxpayer burden from audit compliance, and government expenditures from the audit. With the universe of individual income tax returns in Pakistan from 2012 – 2020, we employ generalized random forests to estimate the individualized causal effects and optimize the distribution of audits with stochastic gradient descent and genetic algorithms. We find that the welfare cost per-dollar of revenue raised can be reduced by between 40% – 57% while collecting even more revenue than under the observed policy.

1 Introduction

Public economics has long been concerned with the relative welfare impacts of alternative policies. This has led to the popularity of cleverly formulated structural welfare models from which we can derive sufficient statistics out of model primitives. Some popular examples of such sufficient statistics include the Marginal Excess Burden (e.g. Eissa et al. (2008); Eissa and Hoynes (2011)), Net Social Benefit (Olken, 2007), and the Benefit-Cost Ratio (e.g. Heckman et al. (2010)). The most recent advancement in this literature, the Marginal Value of Public Funds (MVPF) (e.g. Mayshar (1990); Slemrod and Yitzhaki (1996); Hendren (2016); Hendren and Sprung-Keyser (2020)), extends the work of these previous approaches by formulating a sufficient statistic in terms of empirically estimable causal effects. In many cases this provides a major advantage over sufficient statistics which rely on abstract Hicksian elasticities that may be difficult to identify in practice.

To this point, the MVPF has generally been used for a single purpose: to compare the relative welfare impacts *across* policy domains using *average* causal effects to derive the MVPF itself. That is, economists have aimed to answer the question of "how much should we spend on policy A vs. policy B?" when policy A might be a marginal tax rate reform and policy B might be an expenditure program like school vouchers, unemployment insurance, etc. (e.g. Bergstrom et al. (2024)). Meanwhile, the field of econometrics has recently exploded with advances in methods to estimate heterogeneous causal effects with non-parametric machine learning (e.g. Wager and Athey (2018), Athey et al. (2019)). This potentially opens the door for another important question of policy optimization: "*who* should be treated under a fixed budget to maximize the social welfare impact of the policy itself?"

In this paper, we provide a framework for deriving the optimal distribution of treatment for empirical welfare maximization with the MVPF as the guiding objective function. The framework is developed through the use-case of audit policy in Pakistan. It is well-established that tax evasion rates are high in developing economies for a variety of reasons including (but not limited to) a pervasive informal sector and small budgets for enforcement (Jensen, 2022). Developing economies collect between 10-20% of their GDP in taxes, as opposed to their OECD counterparts which collect around 40% on-average (Besley and Persson, 2014), despite having similar marginal tax rates on income. Economists have posited for some time that a lack of tax enforcement capabilities is a significant barrier to economic development, stated as early as Kaldor (1963), "It is the shortage of resources, not adequate incentives, which limits the pace of economic development." This remains an issue today, as evidenced by a recent IMF policy briefing (Benitez et al., 2023) which claims "...developing countries have made some progress in revenue mobilization during the past decades... however, much more is needed," and by a recent World Bank report (Dom et al., 2022) which states that, "In many developing countries, tax revenues remain far below levels needed to provide citizens with basic services or fund extra spending to minimize the impact of COVID-19."

Given the high rates of tax evasion and the limited audit budgets we observe in developing economies, it becomes especially critical to optimally select the audit population. In the case of Pakistan, 80% of audited firms were found to be evading over a 3-year period where audits were completely randomized among all firms in the country (Farooq, 2024). The hypothesis of this paper becomes clear in light of this empirical fact: if nearly everybody is evading, then determining audit policy based only on the likelihood of evasion is not useful. Rather, audit policy should be determined optimally based on the expected welfare cost of that policy.

Using administrative data from Pakistan which covers the universe of individual income tax returns and audits from 2012 – 2020, we apply new developments in causal machine learning techniques (Wager and Athey (2018), Athey et al. (2019)) which allow us to estimate the causal response functions of individuals to audits using observables. With this mapping, we derive optimal targeting schemes which seek to minimize preference-weighted MVPFs, which is equivalent to minimizing the welfare cost per-dollar of government revenue raised from audits.

We begin our analysis by deriving a simple structural model which identifies the relevant treatment effects that together allow the derivation of MVPFs as a function of the treatment distribution. By estimating these treatment effects at the individual level, and by making some assumptions on the nature of the government's redistribution preferences, these MVPFs serve as sufficient statistics to compare the relative welfare impacts of alternative treatment distributions. We establish three causal effects that are necessary to derive our MVPF. (1) The net present value (NPV) of future tax revenue streams, which accounts for both the uncertain revenue recoup from an audit *and* all other changes in future economic behavior resulting from the audit – e.g. deterrence, future business profits[1], etc. (2) The expected marginal cost to the government from an audit of any given individual. (3) The expected costs accrued to the taxpayer from having to comply with the audit process – e.g. hiring accountants/lawyers, time waste from having to provide documentation to assessors, etc.

Next, we outline our procedure to estimate these individualized treatment effects with causal machine learning methods. Specifically, we leverage a year where, by happenstance, all federal audits in Pakistan were conducted completely at-random. This allows us to apply generalized random forests to estimate individualized treatment effects and apply them to the universe of tax returns in the following year. With these individualized treatment effect predictions, we use versions of stochastic gradient descent and genetic algorithms to derive empirically optimal policies. We consider four policy objectives the government may reasonably seek to pursue: (1a) minimize the social cost per-dollar of revenue raised (i.e. the MVPF) subject to a minimum revenue recoup threshold, (1b) minimize the MVPF subject to both a revenue recoup minimum and a maximum audit expenditure budget, (2a) maximize the NPV of revenue recouped subject to a maximum MVPF threshold, (2b) maximize the NPV of revenue recouped subject to both a maximum MVPF constraint and a maximum audit expenditure constraint. Finally, we compare the social costs and revenues/expenditures of the optimal predicted policies to the observed policy, and conduct several tests to provide evidence for the validity of these counterfactual policy gains.

[1]It has been documented in developing economies that enforcing statutory tax rates may drive some firms to the informal sector as some economic activities will not be profitable under statutory rates.

We find large expected reductions to the welfare cost per-dollar of revenue recoup and equally large expected gains to the revenue recoup amounts themselves under each policy objective from our optimal derived policies relative to the observed audit distribution in 2017. For policy goals 1a and 1b, we predict a reduction in the MVPF (which in this case corresponds to a reduction in the social cost of revenue recouped) by 40% – 57% at no cost to aggregate revenue recoup (in fact, we observe large gains to revenue recouped). For policy goals 2a and 2b, we estimate that revenue recoup can be more than doubled while maintaining the same social cost as under the observed policy.

2 Theory

We begin with a general model which motivates the MVPF framework in the setting of audit policy, keeping in mind earlier and current work on optimal tax enforcement (Keen and Slemrod, 2017) and audit MVPFs (Boning et al., 2023). The purpose of this model is to establish the MVPF as a theoretical benchmark (a sufficient statistic, of sorts[2]) from which we may directly compare the relative efficiencies of audit schemes that target various population subsets, allowing flexibly for a variety of welfare preferences to enter.

2.1 A General Welfare Model in an Audit Setting

We observe individuals indexed $i \in I$ over $t \in T$ periods, who have utility functions over generic consumption, c_{it}, and declared income, y_{it}. Earnings are taxed according to a liability function, $T(y_{it})$, but individuals may choose to evade e_{it} dollars of income taxes such that their consumption in period t is given by $c_{it} = y_{it} - T(y_{it}) + e_{it}$ if individual i is not audited. In order to allow for the dynamic impact of audits on behavior, denote $\alpha_{it} = (a_t, ..., a_{t+k})$ as a vector of audit-year indicators where $a_t = 1$ if the individual was audited in period t. Also denote a penalty function, $\phi(e_{it})$, that represents the additional amount paid by the individual over-and-above their evasion amount if they are audited which is an increasing function of their evasion amount, such that if the individual is caught evading, they will pay $[\phi(e_{it}) + e_{it}]$ to the enforcement agency.

Individuals seek to maximize their expected present discounted value of utility by selecting income y_{it} and an evasion amount e_{it} in each period within their utility function:

$$\max_{y_{it}, e_{it}} \quad E\left[\sum_{t}^{T} \beta^{t-1} \left(c_{it} - \psi(y_{it}) - B_i \cdot a_t\right)\right] \tag{1.1}$$
$$\begin{aligned} s.t. \quad c_{it} &= y_{it} + e_{it} - T(y_{it}) \quad \text{if} \quad a_t = 0 \\ c_{it} &= y_{it} - \phi(e_{it}) - T(y_{it}) \quad \text{if} \quad a_t = 1 \end{aligned}$$

Where per-period utility is assumed to be given by the quasi-linear function, $u(c_{it}, y_{it}) = c_{it} - \psi(y_{it}) - B_i \cdot a_t$. Disutility of labor is given by the function $\psi(y_{it})$, and taxpayer disutility of being audited (also can be interpreted as the marginal taxpayer burden of an audit) is given as B_i. Where the expectation comes from the probability of being audited ($a_t = 1$) in each period, denoted $p_{it}(\boldsymbol{X}_{it}) = Pr\{a_t = 1\}$, and is a function of observable individual i characteristics, $\boldsymbol{X}_{it}$, which may include (but not limited to) past audit history and current/past income tax return information.

Because income, y_{it}, enters additively here, we can focus solely on the individual's evasion decision as the two choices are independent (i.e. the choice of labor earnings is not influenced by the choice of evasion amount). Plugging the constraints into the objective function yields:

$$\max_{e_{it}} \quad \sum_{t} \beta^{t-1} \Big[p_{it}\big(y_{it} - \phi(e_{it}) - T(y_{it}) - \psi(y_{it}) - B_i\big) + (1 - p_{it})\big(y_{it} + e_{it} - T(y_{it}) - \psi(y_{it})\big)\Big] \tag{1.2}$$

[2]The MVPF is, of course, *only* sufficient unless we are willing to make assumptions about social welfare preferences, of which we will explore a variety.

The optimal evasion amount in each period, e^*_{it}, will ignore all terms where e_{it} does not enter, and therefore will solve in each period:

$$\max_{e_{it}} \quad (1-p_{it})e_{it} - p_{it}\phi(e_{it}) \tag{1.3}$$

The solution follows as the general result of first-order conditions for quasi-linear utility:

$$(1-p_{it}) = p_{it}\frac{\partial\phi(e_{it})}{\partial e_{it}} \implies e^*_{it} = \Big(\frac{\partial\phi(e_{it})}{\partial e_{it}}\Big)^{-1}\Big(\frac{1-p_{it}}{p_{it}}\Big) \tag{1.4}$$

Plugging the optimal evasion amount into the objective function yields the expected indirect utility function (i.e. maximum value function of utility) in period t, ex-ante:

$$\begin{aligned} V_i\Big(p_{it}, \phi(e^*_{it}), T(y^*_{it})\Big) &= E\Big[\sum_t \beta^{t-1}\Big(y^*_{it} - \psi(y^*_{it}) - T(y^*_{it}) + e^*_{it} - a_t(e^*_{it} + \phi(e^*_{it}) + B_i)\Big)\Big] \\ &= \sum_t \beta^{t-1} Pr\{\alpha_{it}\}\Big(y^*_{it} - \psi(y^*_{it}) - T(y^*_{it}) + e^*_{it} - p_{it}(e^*_{it} + \phi(e^*_{it}) + B_i)\Big) \end{aligned} \tag{1.5}$$

Where $Pr\{\alpha_{it}\}$ is the probability of a particular sequence of future audits.

2.1.1 Willingness to Pay for an Audit Reallocation

What is the welfare impact of changing the probability of audits at the individual-level? It is helpful to first consider the impact of a general exogenous increase in the audit probability before considering hypothetical budget-neutral reallocations (i.e. increasing the probability of auditing some financed by decreasing the probability of auditing others), as in Boning et al. (2023). Let's consider a uniform increase in the audit probability for type-i individuals in the first period, dp_{i1}. Individuals are willing to pay the value of their compensating variation to avoid an audit: the amount of money from their own income such that their utility is equated between the event of an audit vs. no audit. This is well-established to be equivalent to $\frac{dV_i}{dp_{i1}}$. The individual's indirect utility in this period is given by:

$$V_i = p_{i1}[\underbrace{y^*_{i1} - \psi(y^*_{i1}) - T(y^*_{i1}) - \phi(e^*_{i1}) - B_i + \beta V^1_i}_{\text{Utility if audited}}] + (1-p_{i1})[\underbrace{y^*_{i1} - \psi(y^*_{i1}) - T(y^*_{i1}) + e^*_{i1} + \beta V^0_i}_{\text{Utility if not audited}}] \tag{1.6}$$

Where V^1_i and V^0_i are the net-present-values (NPVs) of utility in subsequent periods if the individual is audited or not in period 1. The envelope conditions tell us that the behavioral responses to the change in audit probability ($\frac{dy^*_{i1}}{dp_{i1}}$ and $\frac{de^*_{i1}}{dp_{i1}}$) will not enter the derivative of the indirect utility function with respect to the audit probability, leaving us with the following:

$$\begin{aligned} \frac{dV_i}{dp_{i1}} &= y^*_{i1} - \psi(y^*_{i1}) - T(y^*_{i1}) - \phi(e^*_{i1}) - B_i + \beta V^1_i - y^*_{i1} + \psi(y^*_{i1}) + T(y^*_{i1}) - e^*_{i1} - \beta V^0_i \\ &= -\phi(e^*_{i1}) - e^*_{i1} - B_i + \beta V^1_i - \beta V^0_i \\ &= -\Big[\underbrace{\phi(e^*_{i1}) + e^*_{i1}}_{\substack{\text{Mechanical (short run)} \\ \text{impact on gov't revenue}}} + \underbrace{\beta(V^0_i - V^1_i)}_{\substack{\text{NPV of long run} \\ \text{gov't revenue}}} + B_i\Big] \end{aligned} \tag{1.7}$$

The first term, $\phi(e^*_{i1}) + e^*_{i1}$, is the immediate tax revenue increase from an audit of individual i, and the second term $\beta(V^0_i - V^1_i)$ represents the NPV of what we broadly refer to as the "deterrence effect" – i.e. the behavioral response to an audit on future tax revenue.[3] Therefore, individual i is willing to pay the sum of this immediate revenue impact (denote as $R^{immediate}$), the NPV of the difference in their future tax revenue streams (denote as R^{future}), and the value of their audit burden (e.g. value of their time filing paperwork, expenses paid to attorneys or accountants, etc.), B_i, to avoid an audit. This outlines the set of causal effects one must estimate for individual i in order to identify their willingness to pay to avoid dp_{i1}:

$$WTP_i = -\frac{dV_i}{dp_{i1}} = R^{immediate}_i + R^{future}_i + B_i \tag{1.8}$$

[3]It is key to note that there are potentially many factors outside of "deterrence" which may impact this value. For example, an individual's business my shut down if they are audited, which may lower income and therefore future tax revenue. In such an event, the "WTP to avoid an audit" falls.

2.1.2 Government Net Costs and the MVPF

The second component of the MVPF as outlined by Hendren (2016), Hendren and Sprung-Keyser (2020), and several others, is the NPV of net program costs to the government. In the case of audit policy, government revenue is affected by 4 sources: (1) collected tax revenue from declared income, (2) evasion amounts, (3) tax revenue recouped from audits, and (4) the marginal costs of audits. Let G_i denote the expected NPV of government revenue from individual i:

$$G_i = \sum_t \beta^{t-1} E\Big[T(y_{it}^*) + a_{it}\big(e_{it}^* + \phi(e_{it}^*) - C_i\big)\Big] \tag{1.9}$$

Where C_i is the marginal cost of auditing person i. Again, the expectation operator is with respect to the probability of auditing person i, where $a_{it} \in \{0, 1\}$ according to this conditional probability.

In a similar manner to the taxpayer's problem, let's expand the first period NPV of government revenue from person i, then assess the impact of a change in audit probability in the first period. Let's first note that we cannot apply the envelope theorem here as the behavioral responses, $\frac{de_{it}^*}{dp_{i1}}$ and $\frac{dy_{it}^*}{dp_{i1}}$, have first-order impacts on government revenue. However, in the first period, the change in audit probability is unknown to the taxpayer. So, by definition, they will have no behavioral response in the first period, though it is worth noting that these responses are the key causal effects driving "deterrence" and are thus part of G^1 and G^0. Differentiating therefore gives us the following, where $G^1 - G^0$ is equivalent to $V^0 - V^1$ as both are simply the NPV of the additional taxes paid under a state of audit vs. a state of no audit (to see this, note that B_i is not a factor outside of period 1, which means that the only relevant component of V and G is R_i under each possible state):

$$G_i = p_{i1}[T(y_{i1}^*) + e_{i1}^* + \phi(e_{i1}^*) - C_i + \beta G^1] + (1 - p_{i1})[T(y_{i1}^*) + \beta G^0] \tag{1.10}$$

$$\begin{aligned}\frac{dG_i}{dp_{i1}} &= T(y_{i1}^*) + p_{i1}T'(y_{i1}^*)\underbrace{\frac{dy_{i1}^*}{dp_{i1}}}_{=0} + e_{i1}^* + p_{i1}\underbrace{\frac{de_{i1}^*}{dp_{i1}}}_{=0} + \phi(e_{i1}^*) + p_{i1}\phi'(e_{i1}^*)\underbrace{\frac{de_{i1}^*}{dp_{i1}}}_{=0} - C_i + \beta G^1 \\ &\quad + T'(y_{i1}^*)\underbrace{\frac{dy_{i1}^*}{dp_{i1}}}_{=0} - T(y_{i1}^*) - p_{i1}T'(y_{i1}^*)\underbrace{\frac{dy_{i1}^*}{dp_{i1}}}_{=0} - \beta G^0 \\ &= \underbrace{\phi(e_{i1}^*) + e_{i1}^*}_{\substack{\text{Mechanical (short run)} \\ \text{impact on gov't revenue}}} + \underbrace{\beta(G^1 - G^0)}_{\substack{\text{NPV of long run} \\ \text{gov't revenue}}} - C_i\end{aligned} \tag{1.11}$$

Let $R_i = R_i^{immediate} + R_i^{future} = \phi(e_{i1}^*) + e_{i1}^* + \beta(G^1 - G^0)$, which is the total impact on government revenue (immediate revenue recoup + long-run "deterrence" effect) from an audit of person i, before costs. The MVPF of any particular audit distribution within the set of all possible audit distributions ($A_p \in \boldsymbol{A}$) is given by [4]:

$$MVPF_{A_p} = \frac{\sum_{i \in A_p}(R_i + B_i)}{\sum_{i \in A_p}(R_i - C_i)} = \frac{\sum_{i \in A_p}(\frac{R_i}{C_i} + \frac{B_i}{C_i})}{\sum_{i \in A_p}(\frac{R_i}{C_i} - 1)} = \frac{\text{Total WTP to Avoid an Audit}}{\text{Gov't revenue recoup} - \text{Gov't costs}} \tag{1.12}$$

2.1.3 Using MVPFs to Compare Policies

Equation 1.12 allows for heterogeneity in behavioral responses and marginal costs across the population to impact the value of the MVPF depending on which taxpayers the government chooses to audit. In principal, if one could identify individualized values of R, C, and B, then one could seek to optimize the MVPF

[4] The numerator here technically represents the WTP to avoid dp_{i1}, which is to say they are willing to pay $\sum_i(R_i + B_i)$ to avoid an increase in their audit probability by dp_{i1}. A possible issue for further examination is, since dp_{i1} is likely not marginal when considering the *selection* of audits (i.e. increasing the probability of an audit from p_{i0} to 1), the numerator may poorly approximate the true WTP to avoid an audit.

by carefully selecting audits. In this case, the MVPF represents the total welfare cost imposed on taxpayers for each \$1 of revenue raised by the government. Thus, a lower MVPF is better: for example, an MVPF of 1.13 would indicate that each \$1 of revenue recouped by the government imposes a \$1.13 welfare cost on taxpayers. Similarly, an MVPF = 1 would indicate a non-distortionary audit policy. Note that negative MVPFs do not indicate better policies as they normally would in cases where program recipients benefit positively from treatment (for example, a policy recovering \$1 in revenue with substantial government costs would yield a negative MVPF), and MVPFs between 0 and 1 are also poor policies as they are only mathematically possible when $\sum_i R_i$ is negative (assuming B_i is always positive), which indicates the audit scheme generates *less* long run revenue to the government than if the government had forgone an audit program entirely. Therefore, the optimal policy would produce an MVPF closest to 1 among the set of policies where $MVPF_{A_p} \geq 1$, and going forward we will only consider this set of policies.

However, a direct comparison of raw MVPFs is not sufficient to make claims about the relative welfare impacts of each alternative policy unless we are willing to assume Utilitarian social welfare preferences[5]. While such an exercise is a useful starting point, economists are typically not willing to make such an assumption as uniform Pareto weights do not account for income inequality aversion or subjective redistribution preferences. Hendren and Sprung-Keyser (2020) – and many others – define the "social marginal utility of income" for person i, denoted η_i, as the marginal increase (decrease) in social welfare when endowing (taxing away from) person i an additional \$1 of income. To briefly illustrate, consider a social welfare function that is a weighted average of individual utilities: $W = \int_i \gamma_i V_i \, di$, where γ_i is individual i's Pareto weight – i.e. the marginal increase in W for each additional util given to person i. If the marginal utility of income from person i is given by $\frac{\partial V_i}{\partial y_i} = \lambda_i$, then the *social* marginal utility of income is how much society values giving that additional \$1 to person i, which will be $\gamma_i \lambda_i = \eta_i$.

In general, when comparing any two policies using their MVPFs in a hypothetical budget-neutral scheme where an increase in spending on policy 1 is financed by an equivalent decrease in spending on policy 2, it is straightforward to show that policy 1 is preferred to policy 2 iff:[6]

$$\bar{\eta}_1 MVPF_1 < \bar{\eta}_2 MVPF_2 \implies \frac{\bar{\eta}_1}{\bar{\eta}_2} < \frac{MVPF_2}{MVPF_1} \tag{1.13}$$

That is, the MVPF of policy 1, weighted by the average social marginal utility of policy 1's beneficiaries, must be less than the same measure for policy 2 in order to conclude that policy 1 is "better" for social welfare as this implies that the preference-weighted welfare cost per-dollar of revenue raised is smaller in policy 1. Therefore, an empirical approach which calculates and compares MVPFs must consider alternative sets of welfare preferences, as MVPFs can only be compared empirically if one assigns a welfare weight to the set of audited individuals. In Section 4.3, we will outline our empirical approach for the consideration of welfare preferences, which amounts to defining a scaling function of tax revenue (and burden) along the income margin such that \$1 of revenue (burden) collected from poorer individuals counts less (more) towards the MVPF than wealthier individuals. The degree of inequality aversion is controlled by a concavity parameter which determines how much less (more) this \$1 of revenue (burden) counts towards the MVPF.

3 Data: Pakistan's Audit Program

We observe the universe of individual tax returns in Pakistan from 2012 – 2020. This includes the full panel of tax returns for approximately 2.8 million individuals, including approximately 31,000 each year who were audited. Our data contain all covariate information within tax returns such as income sources, declared deductions and withholdings, filing dates, etc. We also observe detailed information on audits themselves from 2014 – 2020. This includes a comprehensive list of audited individuals, associated audit dates, and tax recoup amounts.

[5]It's not sufficient because the welfare costs are concentrated on the treated population (i.e. the audited folks). So, we care about not just how much of a cost that is being imposed, but also who we are imposing that cost on.

[6]Note that in the case where the policy is beneficial to recipients (as opposed to here where it is harmful to those treated), this inequality is reversed as larger preference-weighted MVPFs are better.

Regarding costs to the tax agency, our information is less granular, though it remains sufficient (in our opinion) to get decent per-audit cost estimates. We observe tax office-level annual costs along with the individuals who were audited in that year. These data contain fairly detailed information of office-level expenditures: wages paid to assessors and hourly rates, various overhead items, etc. – the limitation being that we cannot directly match expenditures to individual audits. For our first draft of results, we assume the per-audit cost of an individual is equal to the average cost of audits within the tax office they were assigned. To estimate the average cost of audits for a given office $\times$ year, we refer to qualitative accounts for about how much of each budgetary line item could be reasonably attributed to audit. For instance, approximately 1/2 of an auditor's time is spent actually conducting audits, so we attribute the average cost of labor to be 1/2 the annual wages paid to assessors. We also include the costs of computer equipment and office supplies in our estimate. In future drafts, we aim to refine this estimate with further information from domain experts in the Federal Board of Revenue regarding approximate time to complete different audits, which assessors would be assigned to which audit, what overhead items would be employed by different audits, etc.

Taxpayer burden is interpreted as the cost of time and effort (in dollars) imposed on the taxpayer by means of producing paperwork or other documents related to the audit itself, hiring lawyers and/or accountants, or any other monetary expenses associated with complying to the audit itself. This value is not observable in our data, and will likely vary primarily according to how imposing the assessor is on the taxpayer. For now, we translate estimates from a survey conducted in the US by Guyton and Ii (2021) to real Pakistani Rupees (PKR). They estimate an average taxpayer burden on audited individuals of \$3,198, which equates to about 520,085 PKR in 2021. They also estimate a regression akin to $B_i = \alpha + \boldsymbol{\gamma' X_i} + \epsilon$, where $\boldsymbol{X}_i$ is a vector of survey variables. They estimated a coefficient on log(income) to be 0.18, indicating the total burden increases by 0.18% on-average with a 1% increase in income. We extrapolate this estimate across the income distribution to generate (very moderate) heterogeneity in taxpayer burden estimates as a function of income. We also follow the approach taken by Boning et al. (2023), who also refer to this survey. They estimated an average B/C ratio of \$0.50, and chose to hold this constant for all individuals. That is, for an individual with an estimated audit cost of \$1,000, they impose a burden estimate of \$500. The idea is that most of the heterogeneity in taxpayer burden will come from the degree to which the auditor himself imposes it. In our case, the average cost to conduct an audit is XXXXX PKR. Dividing the 520,085 PKR estimate for average burden by this value yields a B/C ratio of VVVVVV. We also keep this ratio constant across the distribution of costs in order to generate most of the heterogeneity in taxpayer burden.

Before 2011, tax enforcement in Pakistan was decentralized: Federal Board of Revenue (FBR) offices determined audit cases independently. However, post 2011, the government began rolling out a centralized auditing program where audits were determined based on a randomized or semi-randomized lottery. Figure 1 depicts a timeline for each of the three types of audit schemes implemented by the FBR over this 8-year period. Specifically, from 2012 – 2014, and again in 2016, audits were conducted completely at-random among the full universe of taxpayers. From 2015 – 2018, excluding 2016, audits were randomized among an eligible sample of taxpayers classified as "likely evaders." The methodology of this classification is proprietary information to the FBR, yet certain covariates are considered "flags" and are published on their website[7]. Finally, from 2018 – 2020, audits were randomized among an eligible sample of taxpayers classified as "noncompliant." Again, the exact methodology for this classification is not available, but noncompliance flags include violations of tax return filing rules – e.g. late returns, improper declarations, etc. In all cases we observe the full list of audited taxpayers as well as those taxpayers that were eligible but not audited.

[7] https://www.fbr.gov.pk/categ/audit-policy/126

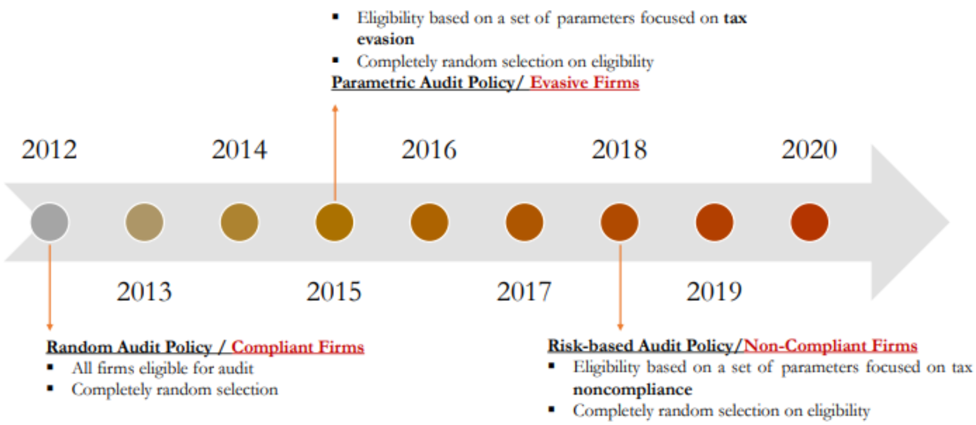

Figure 1: Timeline of FBR's Audit Policy

Notes: Figure comes from Farooq (2024). This depicts the 8 years of natural experiments in audit policy in Pakistan from tax year 2012-2018. During this time Pakistan varied audit eligibility policies between full eligibility, in which all firms are eligible; parametric eligibility, in which only evasive firms are targeted and eligible; and risk-based eligibility, in which only non-compliant firms are targeted and eligible.

4 Machine Learning Conditional Average Treatment Effects and Policy Optimization

This section will discuss our empirical approach to estimating individualized treatment effects and using those individualized treatment effects to derive an optimal policy according to various policy objectives. Section 4.1 will briefly highlight how the randomization of audits in 2016 allows us to create identifying equations for the average impact of an audit on (1) R_i: the NPV of current and future tax revenue, (2) C_i: the marginal costs to the government from the audit, and (3) B_i: the taxpayer burden. Section 4.2 then discusses how we can generalize these identifying equations and apply generalized random forests to estimate individualized taxpayer treatment effects along R, C, and B. Finally, in Section 4.3 we discuss how we apply stochastic gradient descent and genetic algorithms to solve the high dimensional non-linear optimization problems that we propose as policy objectives.

4.1 Identification of MVPF components

The randomized audit schemes from 2012 – 2014 and 2016 provide experimental variation in audit exposure across the distribution of Pakistani taxpayers. Let us first discuss how one should construct and estimate an average treatment effect of audits on R_i. Consider an audit of taxpayer i in the year 2016. If we wanted to estimate the *average* "deterrence" effect, we could estimate a difference-in-differences with the average change in NPV tax revenues from 2017 vs. 2016 among audited vs. non-audited individuals. Recall that the effect on tax revenue in the current period from an audit is simply the revenue recoup: $\phi(e^*_{i1}) + e^*_{i1}$. In subsequent periods, the difference in tax revenue from an audit is the difference in declared tax liability (denote as Y_{it}) over time. This suggests the construction of a variable representing the total NPV of the change in tax revenue over time for individual i, where β is again the exogenous discount rate which we assume in all cases to be 3%: $R^{change}_{i,2016} = \phi(e^*_{i1}) + e^*_{i1} + \beta(Y_{i,2017} - Y_{i,2016}) + \beta^2(Y_{i,2018} - Y_{i,2016}) + ... + \beta^k(Y_{i,2016+k} - Y_{i,2016})$. The average treatment effect of an audit on the NPV of total tax revenue would be identified as τ in the following regression, denoting W_{it} as the indicator for if individual i was audited in period t:

$$R^{change}_{i,2016} = \alpha + \tau W_{it} + \boldsymbol{\gamma'} \boldsymbol{X}_{it} + \varepsilon_{it} \tag{2.1}$$

Where $\boldsymbol{X}_{it}$ is a vector of (optional – in the case of Equation 2.1) individual controls. However, in our application, we do not seek the average impact, τ, but rather we seek the conditional average treatment effect (CATE) function which can be expressed as $\tau(\boldsymbol{X}_{it})$ in the following relaxation of equation 2.1:

$$R_{i,2016}^{change} = \alpha + \tau(\boldsymbol{X}_{it})W_{it} + f(\boldsymbol{X}_{it}) + \varepsilon_{it} \tag{2.2}$$

Where $\tau(\boldsymbol{X}_{it}) = R_i$ from Equation 1.12 and $f(\boldsymbol{X}_{it})$ is now a flexible function of the covariates. The remaining components of Equation 1.12, C_i and B_i, are actually much simpler to construct because they amount to predictions of conditional mean functions under treatment. That is, where $R_i = E[R_{i,t}^{change}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] - E[R_{i,t}^{change}|W_{i,t} = 0, \boldsymbol{X}_{i,t}]$, the equivalent expressions for C_i and B_i are simply:

$$\begin{aligned}
C_i &= E[C_{i,t}^{change}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] - E[C_{i,t}^{change}|W_{i,t} = 0, \boldsymbol{X}_{i,t}] \\
&= E[C_{i,t}^{change}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] - 0 \\
&= E[C_{i,t}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] \\
B_i &= E[B_{i,t}^{change}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] - E[B_{i,t}^{change}|W_{i,t} = 0, \boldsymbol{X}_{i,t}] \\
&= E[B_{i,t}^{change}|W_{i,t} = 1, \boldsymbol{X}_{i,t}] - 0 \\
&= E[B_{i,t}|W_{i,t} = 1, \boldsymbol{X}_{i,t}]
\end{aligned}$$

Where $C_{i,t}^{change}$ and $B_{i,t}^{change}$ are only non-zero in the current period, so the "change" superscript is unnecessary. Again, $\tau(\boldsymbol{X}_{it}) = C_i$ in the case of the cost outcome and $\tau(\boldsymbol{X}_{it}) = B_i$ in the case of the burden outcome. These individualized treatment effect estimates can be used in conjunction with 1.12 to evaluate the social costs of alternative audit distributions.

A few notes are worth mentioning regarding our practical estimation of Equations 2.1 and 2.2. First, because we only observe the distribution of audits in 2016 (of the years where audits were randomized), we estimate CATE functions with only 2016 audits. We consider three years of prior tax returns plus the 2016 return as covariates within $\boldsymbol{X}_{it}$, and we consider three years of future revenue plus 2016 to construct $R_{i,t}^{change}$. This is simply due to the fact that we only observe returns up to 2020, so three is the maximum number of years we observe on the first out-of-bag year (2017). This is so that we may evaluate the quality of CATE function predictions on as much observed data as possible. Nonetheless, our estimates of long-run revenue impacts are understated if deterrence persists longer than three years. Prior work has generally found deterrence to be persistent nearly indefinitely (e.g. DeBacker et al. (2018); Boning et al. (2023)). The consequences of this understatement could be substantive for determinance of the optimal audit allocation in two ways. First, by understating the impact of an audit on long-run tax revenue, the optimal policy will mechanically favor (relatively) auditing individuals with a larger initial recoup prediction over those with possibly larger deterrence effects. If the government's discount rate on future tax revenue is sufficiently low, this could lead to a sub-optimal audit allocation relative to the government's preferences. However, if one is willing to extrapolate the 3-year post-audit deterrence effect to T future periods, as prior work would suggest is reasonable, this problem may easily be remedied. Second, if the rate of "deterrence decay" (i.e. the speed at which declared tax revenue reverts to its' prior level in period 1) differs substantially across taxpayers in unobserved years (i.e. in years $(t + 4)$ through T), this would not be captured in our CATE estimates and any future-year extrapolation of deterrence effects would be problematic. The same prior research has generally not found substantial heterogeneity in the persistence of deterrence across the income distribution, which helps us alleviate that concern.

A second important note is that we restrict our estimating sample to those who both (1) never left the sample between $2016 - 2020$ and (2) were not audited in $2017 - 2020$. The reason for (1) is that the NPV of revenue will be understated in a given year if 0 is imputed for missing year values. We will conduct attrition tests to evaluate the potential that audit status is correlated with attrition, a possibility which seems plausible, though hopefully not substantial. To point (2), we assume the government is selecting the optimal audit allocation in the current year and only in the current year. That is, the government seeks to optimize their objective function by selecting audits only in the current year and not in future years. For this reason, in the training year (2016), we remove individuals who are observed as audited in 2017, 2018, or 2019 as the observed NPV of revenue from these individuals will be affected by treatment in a future year and not the current year, which may introduce a bias on the NPV estimates due to measurement error. In the testing/application year (2017), we remove individuals who are observed as audited in 2018, 2019, or 2020 to compute the observed and predicted MVPFs for the same reason.

4.2 Generalized Random Forests for CATE Estimation

Now that we have established the MVPF as our theoretical benchmark which now can vary within a fixed budget due to heterogeneous behavioral responses and outlined an approach for identification of relevant average behavioral responses, we need a method to map CATE functions such that we can get individualized treatment effect estimates for each component of the MVPF ratio. Here, we turn to the multi-outcome causal forest, a variant of the generalized random forest (GRF) (Athey et al., 2019) proposed by Nie and Wager (2020)[8] which allows for an outcome vector of length J and subsequently J CATE functions.

Let's begin with the data generation processes over the J relevant outcomes:

$$\boldsymbol{Y}_i = f(\boldsymbol{X}_i) + \boldsymbol{\tau}(\boldsymbol{X}_i)W_i + \boldsymbol{\varepsilon}_i$$

Where W_i is the treatment assignment indicator, $f(\boldsymbol{X}_i)$ is a flexible function of the confounders, and $\boldsymbol{\tau}(\boldsymbol{X}_i)$ is the length-J vector of conditional average treatment effects, defined for any given outcome $j \in J$ as the difference in potential outcomes for observation i:

$$\tau(\boldsymbol{X}_i) = E[Y_i(1) - Y_i(0)|\boldsymbol{X_i = x}]$$

The objective of the causal forest (and GRF) is to build a weighting kernel such that we can express the treatment effect of any given observation i as a weighted average of all other observations, where observations with similar treatment effects are weighted more heavily. The forest does this by building B regression trees which recursively split the data along covariates to maximize the squared-difference in predicted ATE's across the resulting splits (denoted C_1 and C_2)[9]:

$$n_{C_1} n_{C_2} (\hat{\tau}_{C_1} - \hat{\tau}_{C_2})^2$$

The result of a singe tree is a set of "neighborhoods" which each contain observations where the predicted ATE is the same (i.e. all observations falling into the same terminal leaf have the same predicted treatment effect for that particular tree). Estimated treatment effects for splitting come from the solution to a local moment condition implied by the identification strategy. In our case with purely random audits, we have experimental variation in treatment assignment. Therefore the conditional exogeneity assumption holds, implying the following population moment for any single outcome:

$$\begin{aligned} E[W_i \varepsilon_i | \boldsymbol{x}_i] &= 0 \\ \implies E[W_i(Y_i - W_i \tau(\boldsymbol{x}_i))|\boldsymbol{x}_i] &= 0 \end{aligned}$$

At the parent node P, all observations are weighted equally because no splits have been made. Therefore the GMM solution is just the sample analog of the above moment where weights = 1, which corresponds to the simple OLS estimate of the average treatment effect (recall that the goal is to estimate the ATE within subgroups – the CATE, but at the parent node the "subgroup" is the full sample, so the CATE = the ATE which is a constant):

$$\begin{aligned} E[W_i Y_i - W_i^2 \tau(\boldsymbol{x}_i))|\boldsymbol{x}_i] &= 0 \\ \implies \hat{\tau}_P &= \frac{\sum_i W_i Y_i}{\sum_i W_i^2} \end{aligned}$$

[8]Their paper actually proposed a method to include multiple treatment arms, however it is easily extended to the case of multiple outcomes so long as the outcomes are on the same scale.

[9]Recursive splitting refers to the continuous process of splitting, then splitting the splits, etc. Regression trees also split "greedily" such that they always select the optimal split from the current node, never re-optimizing or selecting sub-optimal splits even if it could lead to better subsequent splits.

Next, splits are created to maximize heterogeneity in predicted ATE across the resulting nodes (see the objective function above) by iteratively searching across each of the K covariates, $x_k \in \boldsymbol{X}$, and thresholds, c, for the optimal split, where the ATE within resulting splits is estimated via the first-order approximation:

$$\hat{\tau}_c \approx \hat{\tau}_P - \frac{\sum_{i \in c} W_i (Y_i - W_i \hat{\tau}_P)}{\sum_{i \in c} W_i^2}$$

After growing each tree, the CATE function is mapped by considering any point i with covariates $\boldsymbol{X}_i = \boldsymbol{x}$ and then re-estimating the OLS version of the CATE (the one from the parent node) where we add weights to each observation, $\alpha_i(\boldsymbol{x})$, where these weights correspond to the share of total trees any given observation falls in the same terminal leaf as observation i:

$$\hat{\tau}(\boldsymbol{x}) = \frac{\sum_i \alpha_i(\boldsymbol{x}) W_i Y_i}{\sum_i \alpha_i(\boldsymbol{x}) W_i^2}$$

One key note is that in practice, prediction quality is improved when the estimator is so-called "doubly robust" (there are several other benefits associated with this property), which refers to an estimator that relies on both the conditional mean function *and* the propensity score function. Therefore, we use separate forests to estimate this conditional mean function $m(\boldsymbol{x}) = E[Y_i|\boldsymbol{X}_i]$ and propensity score function $e(\boldsymbol{x}) = E[W_i|\boldsymbol{X}_i]$, which we use to de-mean both the outcome and treatment indicators before regressing. I.e. the CATE function in practice is estimated via:

$$\begin{aligned}
\hat{\tau}(\boldsymbol{x}) &= \frac{\sum_i \alpha_i(\boldsymbol{x}) \tilde{W}_i \tilde{Y}_i}{\sum_i \alpha_i(\boldsymbol{x}) \tilde{W}_i^2} \\
&= argmin_\tau \Big\{ \sum_i \alpha_i(\boldsymbol{x}) \Big(Y_i - \hat{m}^{(-i)}(\boldsymbol{X}_i) - \tau(\boldsymbol{X}_i)(W_i - \hat{e}^{(-i)}(\boldsymbol{X}_i)) \Big)^2 \Big\}
\end{aligned}$$

Where $\tilde{Y}_i = Y_i - \hat{m}^{(-i)}(\boldsymbol{X}_i)$, $\tilde{W}_i = W_i - \hat{e}^{(-i)}(\boldsymbol{X}_i)$, and the superscript $(-i)$ simply refers to the fact that the conditional mean and propensity score functions are cross-trained on observations other than i.

The multi-outcome causal forest generalizes this approach to train a single causal forest that can map J CATE functions by incorporating the inner-product of the CATE vector with the de-meaned treatment indicator:

$$\hat{\boldsymbol{\tau}}(\boldsymbol{x}) = argmin_{\boldsymbol{\tau}} \Big\{ \sum_i \alpha_i(\boldsymbol{x}) \Big(\boldsymbol{Y}_i - \hat{\boldsymbol{m}}^{(-i)}(\boldsymbol{X}_i) - \Big\langle \boldsymbol{\tau}(\boldsymbol{X}_i), W_i - \hat{\mathbf{e}}^{(-i)}(\boldsymbol{X}_i) \Big\rangle \Big)^2 \Big\}$$

Intuitively, the above formulation allows Y_i to be vector-valued and for the causal forest to select singular splitting rules that take into account all outcomes at once. Specifically, first note that the objective function of a single-outcome causal forest (maximizing the squared difference in estimated treatment effects across resulting child nodes) can be equivalently expressed as maximizing the sum of the average influence functions between C_1 and C_2:

$$\max_{x_k, C} \quad \tilde{\Delta}(C_1, C_2) = \sum_{j=1}^{2} \frac{1}{n_{C_j}} \left(\sum_{i \in C_j} \rho_{i,P} \right)^2$$

Where $\rho_{i,P}$ is the influence function for observation i on the ATE in a given parent node P, expressed as:

$$\rho_{i,P} = \frac{W_i (Y_i - W_i \hat{\tau}_P)}{\frac{1}{n_P} \sum_{i \in P} W_i}$$

This is a useful reformulation of the objective function because it allows us to concatenate the gradient vector for multiple outcomes into a single objective function for splitting rules:

$$\max_{x_k, C} \quad \tilde{\Delta}(C_1, C_2) = \sum_{j=1}^{J} \sum_{s=1}^{2} \frac{1}{n_{C_s}} \left(\sum_{i \in C_s} \rho_{i,j,P} \right)^2$$

Though estimates of conditional average treatment effect (CATE) functions, $\hat{\tau}^k(X)\, k \in \{1, 2\}$, come from causal forests which leverage the randomized audits of 2016, individual predictions of treatment effects are known to be quite noisy despite their asymptotic consistency (Athey et al., 2019), and the large individual residuals from these predictions are not known to cancel out when aggregating into summary measures such as (in our case) the sum of treatment effect predictions. Fortunately, Robins et al. (1994) showed that the Augmented Inverse Probability-Weighted (AIPW) estimator is asymptotically optimal as it has the lowest variance among all non-parametric estimators. The AIPW estimator of the ATE looks like the following (after some rearranging):

$$\begin{aligned}\tau_{AIPW} &= E\left[\underbrace{\tau(\boldsymbol{X}_i)}_{\text{Initial CATE estimate}} + \underbrace{\left(\frac{W_i - e(\boldsymbol{X}_i)}{e(\boldsymbol{X}_i)[1 - e(\boldsymbol{X}_i)]}\right)}_{\text{Debiasing Weight}} \underbrace{Y_i - m(W_i, \boldsymbol{X}_i)}_{\text{Residual}}\right] \\ &= E[\underbrace{\psi(X)}_{\text{Debiased CATE Estimate / Doubly-Robust Score}}]\end{aligned}$$

Notice that the aggregated measure (in this case the ATE) is simply the average of CATE estimates that have been de-biased via estimates of the propensity score, $e(\boldsymbol{X}_i)$, and the conditional mean, $m(W_i, \boldsymbol{X}_i)$. We do not expect the de-biasing procedure to make a large difference in our case because propensity scores are known and treatment take-up is not optional, nonetheless we employ aggregates of these so-called doubly-robust[10] scores due to their asymptotic optimality.

4.3 Empirical Welfare Maximization

We now turn to the application of individualized CATE estimates to the derivation of a socially optimal audit policy. The MVPF serves as our baseline objective function for a given distribution of audits, A_p:

$$MVPF_{A_p} = \frac{\sum_{i \in A_p}(R_i + B_i)}{\sum_{i \in A_p}(R_i - C_i)}$$

It is key to note that the MVPF is not itself a social welfare function, but instead is used to compare the social cost of alternative audit schemes with empirically estimable causal effects. This implies that a direct optimization of the MVPF, even if we assume Utilitarian welfare preferences, will not make sense in every policy setting, including this one. To illustrate, recall that we previously discussed that the "best" possible MVPF of an audit scheme would be 1, because in our model, the treatment imposes negative "benefits" on those audited and people are willing to pay to *not* receive treatment. This is inverse from usual public policy where benefits are generally assumed to be positive. In this negative benefit setting, the optimal policy according to the MVPF may be to treat a single individual, as this could yield an MVPF very close to 1 (assuming there are no individuals who are costless to audit and do not suffer taxpayer burden from the audit above the revenue recoup amount). However, the socially optimal policy is clearly not to eliminate most tax enforcement as this would ultimately lead to no tax revenue.

So what is the value of the MVPF framework over a traditional social welfare function that can be maximized with respect to the audit distribution? Namely, the MVPF becomes exceptionally useful when constraints are imposed by policy administrators. By minimizing the MVPF subject to constraints on the minimum amount of revenue that should be raised from the audit policy and (optionally) maximum government expenditure constraints, we are able to better answer "what is the minimum social cost the government could have achieved to raise the same revenue had they better selected audits?" By maximizing total revenue subject to a maximum MVPF constraint we can alternatively answer the inverse question: "what is the maximum revenue attainable at the same social cost that was imposed by the existing policy?"

We therefore consider four policy goals the government could pursue:

[10] The property of doubly-robust estimators is that they remain consistent so long as *either* the conditional mean function or the propensity score function is correctly specified.

(1a) Minimize the aggregate MVPF subject to a minimum revenue recoup threshold: $\sum_i R_i \geq \bar{R}$

(1b) Minimize the aggregate MVPF subject to a minimum revenue recoup threshold and a maximum expenditure threshold: $\sum_i R_i \geq \bar{R}$; $\sum_i C_i \leq \bar{C}$

(2a) Maximize revenue recoup subject to a maximum social cost: $MVPF_G \leq \overline{MVPF}$

(2b) Maximize revenue recoup subject to a maximum social cost and a maximum expenditure: $MVPF_G \leq \overline{MVPF}$; $\sum_i C_i \leq \bar{C}$

We follow the "Empirical Welfare Maximization (EWM)" approach defined by Kitagawa and Tetenov (2018). We illustrate the EWM framework by considering policy objectives (1a) and (1b).

The objective of the policy maker is to derive a treatment assignment rule for a "target population" which acts as a new sample drawn from the same joint distribution as the original sampled population. In our case, the target population represents the 2017 realization of the 2016 sample population. The decision rule maps the vector of individual covariates, $\boldsymbol{X}_{it}$, to the binary treatment assignment decision, W_i. Formally, define the full covariate space as χ, and define a decision set, A_p, as some combination of covariate values which map to a treatment decision of $W_i = 1$ if $\{\boldsymbol{X}_{it} \in A_p\}$. Denote the full collection of possible decision sets as $\mathcal{A} = \{A_p \subset \chi\}$, which is the subset of covariate combinations that are both non-random and satisfy our exogenously-imposed constraints (e.g. revenue restrictions, budgetary restrictions, and/or social cost restrictions).

Assume (at first, for simplicity) that we adhere to Utilitarian social welfare preferences. In this case, the social planner's goal is to minimize the expected total welfare cost imposed on taxpayers subject to a minimum revenue threshold and (optionally) a maximum total government expenditure on audits. Which, as we have established, is proportional to the expected MVPF among the audited subsample, where θ represents the joint distribution of outcomes, covariates, and treatment assignment from the sampled population:

$$MVPF(A_p) = \frac{E_\theta\left[\sum_i[(\widetilde{R}_i + \widetilde{B}_i) \cdot 1\{\boldsymbol{X}_{it} \in A_p\}] + \sum_i[(\widetilde{R}_i + \widetilde{B}_i) \cdot 1\{\boldsymbol{X}_{it} \notin A_p\}]\right]}{E_\theta\left[\sum_i[(\widetilde{R}_i - C_i) \cdot 1\{\boldsymbol{X}_{it} \in A_p\}] + \sum_i[(\widetilde{R}_i - C_i) \cdot 1\{\boldsymbol{X}_{it} \notin A_p\}]\right]} \tag{3.1}$$

Equation 3.1 allows us to express the objective function using the entire target population, which will be useful going forward. In practice, outcomes may be weighted to reflect non-Utilitarian social welfare preferences. For example, we can impose different degrees of concavity along the income margin to transform R_i such that revenue collected from poorer individuals is worth less than revenue collected from wealthier individuals (similarly, we can do the inverse of this for B_i, where \$1 of taxpayer burden on a wealthier individual may contribute less than \$1 of social cost if their weight is < 1). In other words, the transformation of R_i should increase positively with income, and the transformation of B_i should decrease (proportionally to the R_i transform) with income. We impose the following transformations:

$$\widetilde{R}_i = \frac{z_i^\alpha}{\bar{z}^\alpha} R_i \quad \widetilde{B}_i = \frac{\bar{z}^\alpha}{z_i^\alpha} B_i$$

Where z_i indicates individual i's income, $\bar{z}$ is the average income of the target sample, and α is an inequality aversion parameter which ranges from 0 (Utilitarian) to ∞ (Rawlsian) social preferences. Figure 2 presents some sample multipliers based on a range of aversion parameters and a hypothetical mean income of 50,000. Going forward, we will omit the $\sim$ for notation simplicity, but in all cases R_i and B_i reflect their appropriately transformed values. This also implies that the constraints reflect welfare-weighted revenue/expense thresholds.

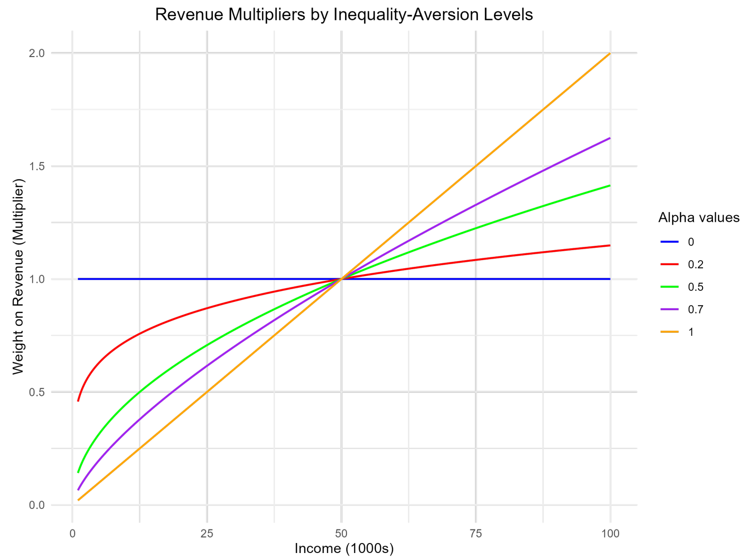

Figure 2: Example Revenue Multipliers

This figure shows some sample multiplier functions over different inequality-aversion levels. At alpha = 0, the welfare weights are equal for all individuals and therefore \$1 of revenue recouped from any individual counts as \$1 in the MVPF. At alpha > 1, \$1 of revenue recouped from someone above (below) the mean of the income distribution counts for more (less) than \$1 when calculating MVPFs.

Equation 3.1 implies that two key causal effect functions should be estimated:

$$\begin{aligned}\tau^1(X) &= E_\theta[(R_i + B_i)|X_i, W_i = 1] - E_\theta[(R_i + B_i)|X_i, W_i = 0] \\ &= m_1^1(x) - m_0^1(x) \\ \tau^2(X) &= E_\theta[(R_i - C_i)|X_i, W_i = 1] - E_\theta[(R_i - C_i)|X_i, W_i = 0] \\ &= m_1^2(x) - m_0^2(x)\end{aligned}$$

Which yields:

$$\begin{aligned}MVPF(A_p) &= \frac{E_\theta\Big[\sum_i \big[m_1^1(x) \cdot 1\{X_i \in A_p\}\big] + \sum_i \big[m_0^1(x) \cdot 1\{X_i \notin A_p\}\big]\Big]}{E_\theta\Big[\sum_i \big[m_1^2(x) \cdot 1\{X_i \in A_p\}\big] + \sum_i \big[m_0^2(x) \cdot 1\{X_i \notin A_p\}\big]\Big]} \\ &= \frac{E_\theta\Big[\sum_i \big[\big(\tau^1(x) - m_0^1(x)\big) \cdot 1\{X_i \in A_p\}\big] + \sum_i \big[m_0^1(x) \cdot 1\{X_i \notin A_p\}\big]\Big]}{E_\theta\Big[\sum_i \big[\big(\tau^2(x) - m_0^2(x)\big) \cdot 1\{X_i \in A_p\}\big] + \sum_i \big[m_0^2(x) \cdot 1\{X_i \notin A_p\}\big]\Big]} \\ &= \frac{E_\theta\Big[\sum_i \tau^1(x) \cdot 1\{X_i \in A_p\} + \sum_i m_0^1(x)\Big]}{E_\theta\Big[\sum_i \tau^2(x) \cdot 1\{X_i \in A_p\} + \sum_i m_0^2(x)\Big]}\end{aligned} \tag{3.2}$$

Because $\sum_i m_0^1(x)$ and $\sum_i m_0^2(x)$ are not functions of the treatment assignment rule, A_p, the policymaker can maximize economic welfare by solving the sample analog of a simplified objective function subject to revenue and budget constraints:

$$\begin{aligned}A^* = \arg\min_{A_p \in \mathcal{A}} \; &\frac{E_\theta\Big[\sum_i \tau^1(x) \cdot 1\{X_i \in A_p\}\Big]}{E_\theta\Big[\sum_i \tau^2(x) \cdot 1\{X_i \in A_p\}\Big]} \\ \text{s.t.} \quad &\sum_{i \in A_p} R_i \geq \bar{R}; \quad \sum_{i \in A_p} C_i \leq \bar{C}\end{aligned} \tag{3.3}$$

Kitagawa and Tetenov (2018) show that $\hat{A}_{EWM}$, the policy derived from solving the sample analog of Equation 3.3, has desirable statistical properties regarding empirical regret. Empirical regret in this setting would be defined as the difference in expected welfare cost from any given derived policy vs. the best feasible policy:

$$E_{pn}\left[MVPF(\hat{A})\right] - MVPF^*_{\mathcal{A}} = E_{pn}\left[MVPF(\hat{A}) - MVPF^*_{\mathcal{A}}\right] \geq 0 \tag{3.4}$$

Where E_{pn} refers to the expectation with respect to different realizations, p, of the random sample, n. Namely, $E_{pn}\left[MVPF(\hat{A}_{EWM})\right]$ converges to $MVPF^*_{\mathcal{A}}$ uniformly with sample size at a rate $O(n^{-1/2})$, which is minimax optimal. Additionally, empirical regret under $\hat{A}_{EWM}$ has a strict upper bound based on n.

4.3.1 Stochastic Gradient Descent Algorithm

Deriving the EWM policies outlined in Equation 3.3 is a high-dimensional combinatorial optimization problem akin to a knapsack problem.[11] We approach the solution with two methods: a version of stochastic gradient descent (SGD) and a genetic algorithm. We employ SGD when deriving policies 1a and 1b, and the genetic algorithm to derive policies 2a and 2b. Referring to Equation 3.1, we can express the objective function with the full sampled population, where the treatment assignment indicator, W_i can be viewed as a parameter/weight:

$$MVPF(A_p) = \frac{\sum_{i \in A_p} W_i(R_i + B_i)}{\sum_{i \in A_p} W_i(R_i - C_i)} \tag{3.5}$$

The SGD algorithm employs a sequence of weighting-and-updating where treatment is randomly assigned (among the set of policies that satisfy the constraints) in an initial pass, and an initial value of the objective function is then computed: $MVPF(\boldsymbol{W}^0)$.

Next, we compute an approximation of the gradient vector of the objective function with respect to a random subset (for computational efficiency) of units in the full sample, $s \in N$:

$$\nabla MVPF^s|_{\boldsymbol{W}^0} = \begin{bmatrix} \frac{\partial MVPF^s}{\partial W_0} \\ \frac{\partial MVPF^s}{\partial W_1} \\ \vdots \\ \frac{\partial MVPF^s}{\partial W_s} \end{bmatrix} \tag{3.6a}$$

$$\frac{\partial MVPF^s}{\partial W_i} = \frac{(R_i + B_i)\sum_i [W_i(R_i - C_i)] - (R_i - C_i)\sum_i W_i(R_i + B_i)}{\left[\sum_i W_i(R_i - C_i)\right]^2} \tag{3.6b}$$

Finally, treatment assignment is updated according to the rule:

$$\boldsymbol{W}^1 = \boldsymbol{W}^0 - \nabla MVPF^s|_{\boldsymbol{W}^0} \tag{3.7}$$

The traditional use of SGD employs a learning rate parameter to determine the step size: how large of an adjustment to make upon the prior step policy. In our case, the weights may only take on values of 1 (treat) or 0 (don't treat), so the step size is better defined as the share of the population used to estimate the gradient vector given that all units within the gradient vector are adjusted according to their marginal effects. That is, if the partial derivative of the $MVPF$ function with respect to observation i is negative, the treatment assignment will be switched on (if it is not already), and treatment will conversely be switched off if the partial derivative is positive. This approach increases convergence speed at the cost of some accuracy in each pass, though we generally use a small subsample (50-100 units) to estimate the gradient vector in each pass, yielding an implicit "step size" which is effectively quite small.

[11] This version of the knapsack problem is a non-linear and non-convex optimization problem in the class of NP-Hard, which means there is no general-purpose algorithm that can guarantee a global minimum solution or, more specifically, there is no algorithm that can solve the problem in polynomial time. Therefore we seek an adequate local minimum that, under a certain degree of (non)convexity, will be near the global minimum.

Our use of SGD guarantees a local minimum with sufficient iterations (epochs). However, our objective function is non-linear and non-convex, making integer programming infeasible and convergence to a global minimum not guaranteed. At this point we are exploring alternative algorithms for mixed integer non-linear optimization that can handle high dimensionality and may guarantee convergence to global minima/maxima.[12]

4.3.2 Genetic Algorithm

The genetic algorithm (GA) is a gradient-free approach to local optimization which serves as an "intelligent" probabilistic search algorithm within the simplex of valid policies. GAs have been used with success in a variety of combinatorial optimization settings (Reeves, 1993), including multidimensional knapsack problems (Chu and Beasley, 1998). They aim to simulate the natural processes of "survival of the fittest" by initializing a binary encoding of random treatment assignment policies and calculating the "fitness" of each policy (which in our case will be the total NPV of revenue). For initial policies which are invalid (i.e. the MVPF exceeds the observed MVPF of the real-world policy), fitness is heavily penalized by the degree to which the policy exceeds the maximum allowable social cost. We initialize 200 policies in the first generation, taking the top 5% (top 10) policies according to fitness and induce "crossover" between policies. This crossover procedure generates another set of 200 policies where random sequences within parent policies are exchanged to create children policies. Additionally, a small random sequence within each child policy is "mutated" (i.e. randomly changed) to induce additional exploration. Fitness is again evaluated and the process repeats until a steady state is reached or until a pre-specified number of generations have been created.

Because the set of valid policies is small relative to the vast set of possible policies, it is unlikely that the initial generation of policies will contain any valid policies. To increase convergence speed and leverage the domain knowledge of the problem, we employ a heuristic which randomly transforms 10 of the first 200 policies into valid solutions by re-assigning treatment to units sequentially and randomly until the resulting policy satisfies the MVPF and (optionally) the cost constraint. Thus, the algorithm begins with at least 10 policies in the valid space from which it may spawn the next generation of policies.

4.3.3 Optimal Decision Trees

* this section is filler for now. May include the Athey and Wager (2021) stuff as an add-on later *

Our version of the EWM approach to policy derivation aims to as much information on CATEs as possible to define treatment assignment rules. As a result, the profile of treated individuals (for example, the minimums/maximums/averages of covariate values in the treated group) may act as an interpretable version of a treatment assignment rule that is otherwise based on a complex machine learning algorithm. In practice, policymakers may prefer simpler and more interpretable treatment rules if there are still welfare gains to be made from them, as they are more transparent than black-box CATE predictions and in some cases may be faster to implement. Athey and Wager (2021) recognize this, and extend the EWM method so that one may create shallow-depth decision trees based on the doubly-robust scores used to calculate aggregates of treatment effect predictions.

5 Analysis of Socially Optimal Audit Policies

This section presents the results of the optimal policy derivation exercises. Section 5.1 begins by comparing the MVPFs, revenues, and costs of the observed 2017 audit policy allocation vs. the predicted values

[12]Note to self: I considered linearizing the objective function with a Taylor series. This provides a linear approximation of the objective function around a given treatment distribution, which would in-turn create a linear programming exercise. However, this is still not computationally tractable because (1) we are restricted to small changes in the distribution and then re-approximating, creating a large sequence of mini-LP problems and (2) the correct small change depends on the full set of possible changes, which is a ton. So I don't think this would hold up in a high-dim setting. This just feels like gradient descent with extra steps anyway. There are other mixed-integer NLP solvers out there worth exploring though, as this is not a new branch of math.

from the optimal policies. This allows us to evaluate the potential gains from targeting according to predicted social cost. We also present descriptions of the treatment group under the optimal policy vs. the observed policy so that we may evaluate how the average audited individual differs under each alternative allocation.

Section 5.2 then presents three evaluation exercises which attempt to provide evidence that the derived policy predictions are representative of the counterfactual MVPFs/revenues/costs had the government implemented them. First, we estimate rank-average treatment effects (RATEs) following Yadlowsky et al. (2023) for revenue to test out-of-bag prediction quality. Second, we apply our trained generalized random forest to the observed audit distribution: evaluating how well the model predicts the MVPF/revenue/costs on the actual out-of-bag observed policy. Finally, we calculate MVPFs on the "overlap sample" – the set of individuals selected for audit under the optimal policy who also happened to be audited by chance in the observed policy. If the GRF can identify heterogeneity effectively, this value should be lower than the observed MVPF.

5.1 Comparing the Optimal Policy vs. the Observed Policy

5.1.1 Optimal Policies vs. Observed Policy Comparison

Tables 1 – 4 display a comparison of MVPFs, Revenues, Costs, and the size of the treated sample in 2017 under the observed policy vs. the optimal policy for each in 1a, 1b, 2a, and 2b, as well as for each set of welfare preferences. These MVPFs must be compared with the welfare weights of the observed policy in mind. That is, the best policy among the class of policies where $\alpha = 0$ is not comparable to the best policy among the class of policies where $\alpha = 0.2$, it is only comparable to policies in the class where $\alpha = 0$. Therefore, the MVPF of the observed policy is only directly comparable to the other policies in the class that use the same welfare weights as were implicit in the observed policy. The decision of which welfare weights the government *should* use is a normative one, though one may seek to back-out the weights which rationalize a given policy as optimal (see, for example, Bargain et al. (2014); Bergstrom and Dodds (2024)), and if these weights differ from any reasonable set of weights, one may conclude the policy itself is sub-optimal. For our purposes, if we assume that the actual redistribution preferences of the government are well approximated by some value of α, then the difference in MVPFs between the observed policy and the best policy under a given value of α represents a real welfare gain.

Please note that, for ease of reading, we convert all currency from 2016 Pakistani Rupees (PKR) to current 2024 US Dollars. These are therefore inflation-adjusted but are *not* scaled to reflect relative purchasing power. All prior estimation (i.e. GRFs, optimization algorithms, etc.) was conducted in the original currency, these tables simply reflect converted values of those results. Also note that in all cases, the constraint we impose on R, C, or the MVPF when deriving the optimized policy is the observed value in 2017. That is, in cases with a minimum revenue threshold constraint, the optimal policy must raise at least (in expectation) the \$3,807,806.46 that was observed as the NPV of government revenue gain from the real-world policy. In cases with a maximum expenditure threshold, the optimal policy can spend no more (in expectation) than the \$2,535,803.13 that the FBR spent in 2017. And in cases with a maximum MVPF constraint, the total welfare cost per-dollar of revenue raised cannot exceed \$4.35.[13]

Table 1: Policy 1a: MVPF Minimization with Revenue Constraint

	Redistribution Preferences		
Metric	Observed.Policy	alpha = 0	alpha = 02
MVPF	4.35	2.63	1.89
Total Revenue	3,807,806.46	5,481,337.21	7,689,131.21
Government Expenditure	2,535,803.13	2,558,963.17	2,716,315.88
Number Treated	2993.00	4119.00	3634.00

[13] For alpha > 0 you are reporting preference-weighted revenues. For these tables you should report unweighted revenues

Table 2: Policy 1b: MVPF Minimization with Revenue and Cost Constraint

	Redistribution Preferences		
Metric	Observed.Policy	alpha = 0	alpha = 02
MVPF	4.35	2.78	2.16
Total Revenue	3,807,806.46	4,958,227.53	6,219,685.36
Government Expenditure	2,535,803.13	2,535,112.10	2,528,896.97
Number Treated	2993.00	4237.00	2786.00

Table 3: Policy 2a: Revenue Maximization with MVPF Constraint

	Redistribution Preferences		
Metric	Observed.Policy	alpha = 0	alpha = 02
MVPF	4.35	4.32	4.34
Total Revenue	3,807,806.46	7,713,256.12	8,312,887.55
Government Expenditure	2,535,803.13	5,101,223.66	5,509,855.14
Number Treated	2993.00	4159.00	3998.00

Table 4: Policy 2b: Revenue Maximization with MVPF and Cost Constraint

	Redistribution Preferences		
Metric	Observed.Policy	alpha = 0	alpha = 02
MVPF	4.35	2.38	2.17
Total Revenue	3,807,806.46	5,648,995.23	6,215,963.89
Government Expenditure	2,535,803.13	2,534,791.69	2,532,119.54
Number Treated	2993.00	4096.00	3997.00

Notes: These tables compare the predicted MVPFs, NPVs of Revenue, and Gov't Costs from the optimized policies vs. the observed MVPF/Revenue/Cost from the actual 2017 audit scheme. We optimize (for now) over two alpha levels. Alpha = 0 corresponds to Utilitarian welfare weights, Alpha = 0.2 corresponds to a moderate degree of inequality aversion. Please note that for this first set of results, we run only 1,000 iterations in SGD and 1,000 GA generations. Therefore these policies may not have all fully converged yet – especially for the GA which tends to have slower convergence in general – and there could be moderate inconsistencies.

For all policy goals (1a – 2b) and all α levels, we find large expected gains to the objective function by following the optimal policy. In Policy 1a (MVPF minimization with minimum revenue constraint), we find that the optimal policy expects to reduce the welfare cost per-dollar of revenue raised by $\approx 40\%$ for Utilitarian weights or by $\approx 57\%$ for $\alpha = 0.2$ weights, while raising much greater revenue than observed. In Policy 2b (MVPF minimization with minimum revenue *and* maximum expenditure constraints), the optimal policy still expects to reduce the welfare cost per-dollar of revenue raised by $\approx 36\%$ for Utilitarian weights or by $\approx 50\%$ for $\alpha = 0.2$ weights, while still exceeding the revenue threshold by a wide margin.

For Policy 2a (Revenue maximization with maximum MVPF constraint), the optimal policy expects to more that double the revenue raised from the observed policy while retaining the same social cost. And for Policy 2b (Revenue maximization with maximum MVPF *and* maximum expenditure constraints), the optimal policy expects to increase revenue by $\approx 50\%$, with only the expenditure constraint acting as binding as the social cost of this policy is well below the observed social cost.

5.1.2 Descriptions of Optimal Treatment Groups

Tables 5 –8 compare the treated group under the observed policy to the treated groups under the optimal policies by displaying the covariate means of treated individuals under each policy for the top 15 variables according to VIF (see Figure 5). The first column displays the covariate means of the treated group within the observed audit sample. The remaining columns display the *difference* in the average of this variable for the optimal treatment group vs. the observed group, along with stars to indicate if this difference is statistically significant. This exercise aims to compare the profiles of the average audited individual in 2017 under the FBR's policy vs. the average audited individual under the optimal policy.

Table 5: Policy 1a: Group Covariate Means

	Observed Policy	Optimal Policy - Observed Policy	
Covariate	Covariate.Mean	alpha = 0	alpha = 02
Gross Business Revenue (L.1)	2236.41778	22342.67***	45176.38***
Income Loss from Business (L.0)	63.63169	2897.91***	2867.38***
Business Income (L.0)	58.37183	99628.98	99528.49
Business Profits (L.0)	76.42958	3543.27***	3639.23***
Turnover (L.1)	1492.17760	19006.18***	33584.07***
Gross Business Revenue (L.2)	2056.16840	29005.19***	48320.75***
Net Business Revenue (L.1)	2188.82904	24126.29***	46185.56***
Business Cost of Sales (L.1)	2344.18130	22114.37***	48006.41***
Salaries Paid (L.0)	42.03999	2641.07***	2797.93***
Net Business Purchases (L.1)	2267.69937	18766.87***	42420.04***
Business Cost of Sales (L.0)	1684.01971	27409.44***	54126.34***
Turnover (L.0)	1061.91774	23268.28***	37783.64***
Net Business Revenue (L.2)	1982.63682	30209.98***	46865.15***
Gross Business Revenue (L.0)	1556.09596	28180.67***	51235.98***
Turnover (L.2)	1398.17230	23220.53***	35762.02***

Table 6: Policy 1b: Group Covariate Means

	Observed Policy	Optimal Policy - Observed Policy	
Covariate	Covariate.Mean	alpha = 0	alpha = 02
Gross Business Revenue (L.1)	2236.41778	23456.78***	45789.12***
Income Loss from Business (L.0)	63.63169	3102.45***	3224.56***
Business Income (L.0)	58.37183	105678.90	106789.01
Business Profits (L.0)	76.42958	3789.01***	3920.14***
Turnover (L.1)	1492.17760	20123.45***	34987.89***
Gross Business Revenue (L.2)	2056.16840	30123.67***	49567.89***
Net Business Revenue (L.1)	2188.82904	25345.78***	47123.45***
Business Cost of Sales (L.1)	2344.18130	23123.45***	49123.56***
Salaries Paid (L.0)	42.03999	2789.01***	2910.12***
Net Business Purchases (L.1)	2267.69937	19789.23***	43987.45***
Business Cost of Sales (L.0)	1684.01971	28945.12***	55345.67***
Turnover (L.0)	1061.91774	24567.89***	38987.12***
Net Business Revenue (L.2)	1982.63682	31567.45***	48123.67***
Gross Business Revenue (L.0)	1556.09596	29567.89***	52567.01***
Turnover (L.2)	1398.17230	24567.78***	36987.89***

Table 7: Policy 2a: Group Covariate Means

	Observed Policy	Optimal Policy - Observed Policy	
Covariate	Covariate.Mean	alpha = 0	alpha = 02
Gross Business Revenue (L.1)	2236.41778	24567.89***	46345.01***
Income Loss from Business (L.0)	63.63169	3250.67***	3376.89***
Business Income (L.0)	58.37183	108456.78	109678.90
Business Profits (L.0)	76.42958	4012.56***	4123.78***
Turnover (L.1)	1492.17760	21567.89***	36234.56***
Gross Business Revenue (L.2)	2056.16840	31567.45***	50987.23***
Net Business Revenue (L.1)	2188.82904	26789.01***	48567.12***
Business Cost of Sales (L.1)	2344.18130	24567.78***	50345.89***
Salaries Paid (L.0)	42.03999	2890.12***	3012.34***
Net Business Purchases (L.1)	2267.69937	21012.45***	45234.78***
Business Cost of Sales (L.0)	1684.01971	30234.56***	56789.01***
Turnover (L.0)	1061.91774	25987.23***	40234.56***
Net Business Revenue (L.2)	1982.63682	32890.12***	49567.45***
Gross Business Revenue (L.0)	1556.09596	30987.34***	53987.67***
Turnover (L.2)	1398.17230	25987.45***	38234.89***

Table 8: Policy 2b: Group Covariate Means

	Observed Policy	Optimal Policy - Observed Policy	
Covariate	Covariate.Mean	alpha = 0	alpha = 02
Gross Business Revenue (L.1)	2236.41778	25678.90***	47234.12***
Income Loss from Business (L.0)	63.63169	3390.34***	3523.45***
Business Income (L.0)	58.37183	112456.78	113678.90
Business Profits (L.0)	76.42958	4234.78***	4356.90***
Turnover (L.1)	1492.17760	22978.56***	37456.23***
Gross Business Revenue (L.2)	2056.16840	32890.45***	51987.34***
Net Business Revenue (L.1)	2188.82904	27890.12***	49987.23***
Business Cost of Sales (L.1)	2344.18130	25987.34***	51789.01***
Salaries Paid (L.0)	42.03999	3012.45***	3145.56***
Net Business Purchases (L.1)	2267.69937	22456.89***	46345.23***
Business Cost of Sales (L.0)	1684.01971	31567.78***	58123.90***
Turnover (L.0)	1061.91774	27234.45***	41567.89***
Net Business Revenue (L.2)	1982.63682	34123.78***	50987.56***
Gross Business Revenue (L.0)	1556.09596	32234.56***	55345.12***
Turnover (L.2)	1398.17230	27234.67***	39567.89***

Notes: These tables compare the profiles of the average audited individual in 2017 vs. the average audited individual under the optimal policies for the top 15 most "important" covariates (according to VIF score) for predicting revenue, cost, and burden treatment effects. Column (1) displays the covariate means for these top 15 variables under the observed policy. Columns (2) and (3) display the differences-in-means between the optimal treatment group vs. the observed treatment group. Variables ending in L.0 represent the value of that tax return variable in the current year, L.1 represents a 1-year lag, and L.2 represents a 2-year lag. We consider only up to 2-year lags in our prediction exercise.

Each policy paints a very similar picture regarding the optimal treatment group: audit richer business owners who claim large deductions. For reference, Pakistan did not audit salary workers during these years

– a fact unknown to the causal forest algorithm but was deduced through the patterns in the data. In all cases the average audited individual under the optimal policy has much larger values of business revenue in the current year and the year prior; consistently claims larger deductions for turnover, costs of sales, and salaries; and often declares losses.

5.2 Evaluation of the Optimal Policy Predictions

First, we train smaller causal forests within the randomization year 2016 only: one which serves as a "test" forest trained on a random 70% subset of taxpayers in 2016, and one which serves as an "evaluation" forest trained on the remaining 30%. The evaluation forest serves as a "source of truth" for the computation of individualized treatment effect functions for the in-sample 30% it was trained on. This allows us to construct rank-weighted average treatment effect (RATE) curves following Yadlowsky et al. (2023) on treatment effect predictions from the test forest applied on the out-of-bag 30%. If the causal forest is able to identify heterogeneity, there will be noticeable gains in total revenues from ranking and sequentially treating units by predicted treatment effect over-and-above a random allocation of treatment. Second, we apply the full causal forest, trained on all 2016 observations, to all observations in 2017. With these out-of-bag predictions, we can calculate a predicted MVPF from the observed policy distribution and contrast this with the actual MVPF of the observed policy distribution. If the causal forest is able to identify the heterogeneity accurately, the predicted MVPF will reflect the observed MVPF. Finally, we leverage the "overlap" individuals audited under both the observed policy and the optimal policy. As discussed in the prior subsection, the observed MVPF of this overlap sample may proxy for the counterfactual MVPF of the full optimal sample.

5.2.1 RATE Estimates

The rank-weighted average treatment effect (RATE) serves as a method for evaluating out-of-bag prediction quality of the causal forest. It comprises of two components inspired by the Receiver Operating Characteristic (ROC) from the machine learning classification literature. First, we compute a "Targeting Operator Characteristic (TOC)" – a function representing the gain over the ATE from ranking units by their CATE prediction and treating sequentially. The TOC function can be expressed as the following:

$$TOC(q) = \underbrace{E\left[Y_i(1) - Y_i(0) | \tau(X_i) \geq F^{-1}_{\tau(X_i)}(1-q)\right]}_{\text{Subsample ATE}} - \underbrace{E\left[Y_i(1) - Y_i(0)\right]}_{\text{ATE}}$$

Where q represents a given quantile of the CATE distribution ($F_{\tau(X_i)}$). To illustrate, at $q = 0.1$, the TOC value will be the difference in the ATE from treating the top 10% of the CATE distribution vs. randomly allocating treatment (i.e. the ATE of the full target sample). The second component of the RATE is the total cumulative gain from ranking known as the "Area Under the TOC (AUTOC)," which is simply calculated as $AUTOC = \int_0^1 TOC(q)\,dq$. The best ranking algorithm would maximize this value among all possible ranking algorithms, but any algorithm which is able to capture heterogeneity will at least yield a positive value for this metric because it will generate gains from ranking.

As a first exercise to understand how well the causal forest predicts on out-of-bag data, we train two auxiliary causal forests using 2016 data only. The first we call a "test forest," and it is trained on a random 70% subsample of the data. We call the second an "evaluation forest," and it is trained on the remaining 30%. The idea is that the evaluation forest generates fitted values of (doubly robust) CATEs for the in-sample 30%, which can be aggregated efficiently into a consistent estimate of the ATE in that sample. We then apply the test forest to that sample, ranking the out-of-bag predictions and counterfactually treating sequentially. Figure 3 displays the resulting TOC for revenue.

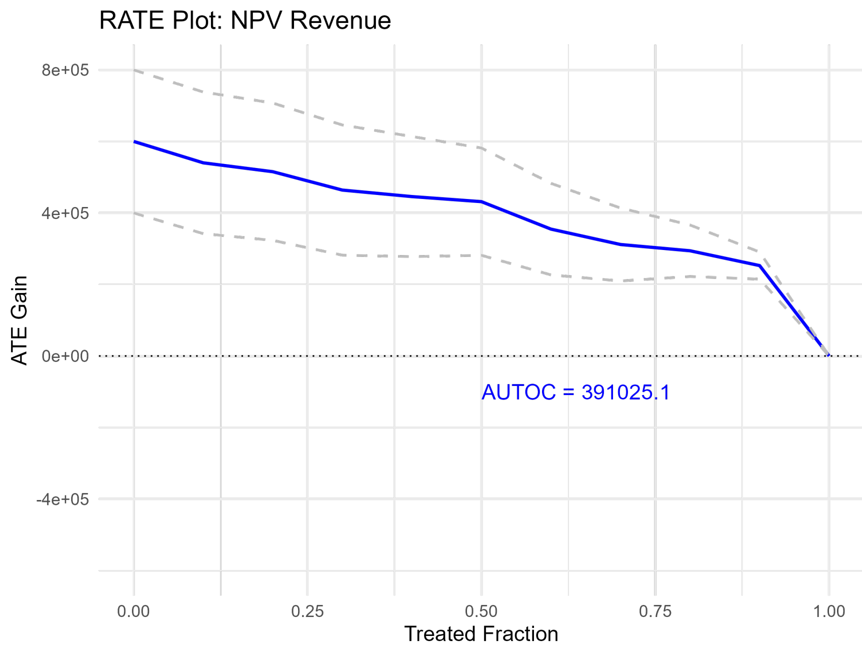

Figure 3: RATE: NPV Revenue

Notes: This figure plots the results of the RATE prediction exercise. It plots the difference between the ATEs of individuals at different deciles of the CATE distribution in the evaluation sample as ranked by the test forest vs. the ATE of the evaluation sample as estimated by the evaluation forest.

The figure suggests that the causal forest does a very good job at identifying heterogeneity in Revenue within the evaluation data. To the extent that the CATE estimates from the evaluation forest serve as a source-of-truth for revenue effects, the test forest generates consistent gains to targeting over the entirety of the treatment effect distribution. That is, there is evidence here that the causal forest is successfully able to rank individuals by their individualized CATEs such that this ranking is reflective of the true ranking.

5.2.2 Predicted MVPF vs. Observed MVPF in 2017 Sample

As a second exercise to evaluate the causal forest prediction quality on the out-of-bag data, Table 9 displays the actual observed MVPF from the audit scheme in 2017 contrasted with the predicted values of this policy from the causal forest:

Table 9: Comparison of Observed and GRF Predicted Values

Metric	Observed	GRF.Predicted	Difference	p.value
MVPF	4.35	4.22	0.13	0.82
Sum of Revenue	3,807,806.46	3,908,563.21	-100,756.75	0.78
Sum of Cost	2,535,803.13	2,559,612.69	23,809.56	0.89
Sum of Burden	1,725,408.03	1,791,728.88	-66,320.85	0.74

Notes: This table tests the causal forest's prediction quality on the observed 2017 treatment distribution. That is, if we are to trust the MVPF/Revenue/Cost estimates from the optimal policies, we need evidence that the causal forest is generally a good predictor of these values on out-of-bag data. Here, we apply the causal forest to the observed distribution (from which we also observe audit outcomes) and conduct two-sample t-tests to evaluate if the prediction error is statistically significant.

This table helps us evaluate the significance of out-of-bag prediction errors on the sample of audits from which we observe outcomes (i.e. the true treatment distribution). In all cases, the prediction error is relatively small, and it is never statistically significant according to a two-sample t-test.

5.2.3 MVPF of Overlap Samples

*** Ignore this section for now. These are not updated – still assessing if I think these are valuable ***

Figure 4 displays the MVPFs of the relative "overlap samples" for policy 1a. They include the observed MVPFs among individuals in the observed policy, the optimal policy, and the individuals in the observed policy who are *not* in the optimal policy. They also include the predicted MVPFs among individuals in the optimal policy but *not* the observed policy, and the entire optimal policy. We include similar Venn diagrams for policies 1b, 2a, and 2b in the appendix.

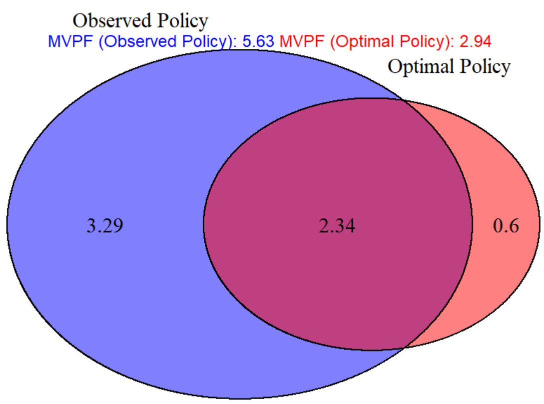

(a) Alpha = 0

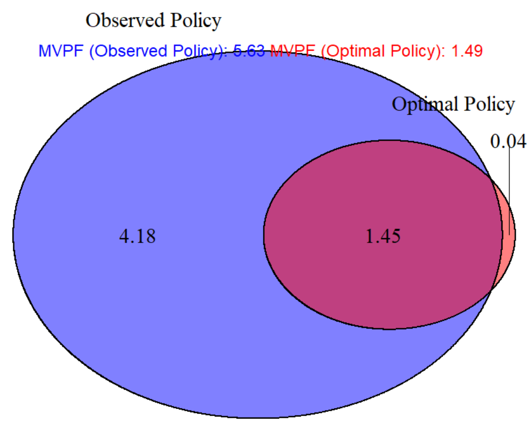

(b) Alpha = 0.2

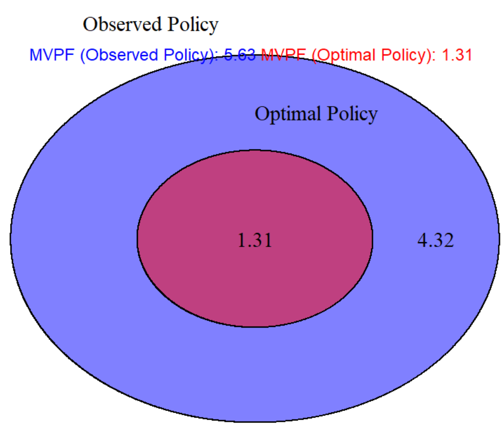

(c) Alpha = 0.5

Figure 4: Venn Diagrams for Policy 1a (MVPF Minimization with Revenue Constraint)

6 Conclusion: How can Governments Incorporate this Framework?

This paper provides a framework for governments to leverage machine learning and welfare theory to derive socially optimal policy allocations. We would like to close with a brief discussion on how this framework may be rolled out by governments in a number of policy settings. Beginning with the case in question, audit policy in Pakistan, there are a few things to keep in mind. First, audit policy provides a convenient setting for machine learning because for most relevant dimensions of the MVPF, the treatment effect predictions are equivalent to conditional mean function predictions under treatment. For example, the "treatment effect" on government costs from auditing any given individual is equivalent to $E[C_i|W_i = 1, \boldsymbol{X}_i]$ because we know that $E[C_i|W_i = 0, \boldsymbol{X}_i] = 0$ (i.e. it is costless to *not* audit someone). The same can be said for taxpayer burden and the revenue recoup in the current period. The only "treatment effect" which relies on randomization to identify is the aspect of revenue which comes from deterrence. Therefore, this is the only aspect of the optimal policy which cannot be estimated with a general-purpose predictive algorithm.

The first solution, which requires further research to confirm, is that deterrence effect predictions may be extrapolated to future years if they are either insignificantly small or relatively constant with time.[14] In this event, a general-purpose predictive algorithm may be applied to the remaining MVPF components and the socially optimal policy may be derived. In this paper, we leveraged the purely-random audit year of 2016 to identify a population-representative distribution of deterrence effects (and other MVPF components). In reality, governments (including Pakistan in nearly every other year) choose to randomize audits among selected subsets of individuals. The optimal policy may still be derived in these settings so long as the government commits to audits within the subset among whom treatment is randomized. So long as the government is able to estimate credible deterrence effects among the sample of individuals from whom we wish to allocate audits optimally, then this framework may be applied. Notably, the same can be said for any policy we wish to allocate optimally: so long as the government can identify treatment effects (or conditional means under treatment when one can credibly assume the conditional mean without treatment is constant) for the population of eligible individuals, this framework may be applied to derive the socially optimal allocation.

References

Alm, J. (2012, February). Measuring, explaining, and controlling tax evasion: lessons from theory, experiments, and field studies. *International Tax and Public Finance 19*(1), 54–77.

Athey, S., J. Tibshirani, and S. Wager (2019, April). Generalized random forests. *The Annals of Statistics 47*(2), 1148–1178. Publisher: Institute of Mathematical Statistics.

Athey, S. and S. Wager (2021). Policy Learning With Observational Data. *Econometrica 89*(1), 133–161. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15732.

Atkinson, A. B. (1970, September). On the measurement of inequality. *Journal of Economic Theory 2*(3), 244–263.

Bargain, O., M. Dolls, D. Neumann, A. Peichl, and S. Siegloch (2014). Comparing inequality aversion across countries when labor supply responses differ. *International Tax and Public Finance 21*(5), 845–873. Publisher: Springer & International Institute of Public Finance.

Benitez, J. C., M. Mansour, M. Pecho, and C. Vellutini (2023, September). Building Tax Capacity in Developing Countries.

Bergstrom, K. and W. Dodds (2024, June). A General Theory of Inverse Welfare Functions. Working Paper 2408, Tulane University, Department of Economics. https://ideas.repec.org/p/tul/wpaper/2408.html.

Bergstrom, K., W. Dodds, and J. Rios (2024, May). Optimal Policy Reforms.

Besley, T. and T. Persson (2014, November). Why Do Developing Countries Tax So Little? *Journal of Economic Perspectives 28*(4), 99–120.

[14] Note that this does not assume there is no decay of the deterrence effect *over* time, it simply means that this rate of decay not change with time.

Boning, W. C., N. Hendren, B. Sprung-Keyser, and E. Stuart (2023, June). A Welfare Analysis of Tax Audits Across the Income Distribution.

Chu, P. and J. Beasley (1998, June). A Genetic Algorithm for the Multidimensional Knapsack Problem. *Journal of Heuristics 4*(1), 63–86.

DeBacker, J., B. T. Heim, A. Tran, and A. Yuskavage (2018, February). Once Bitten, Twice Shy? The Lasting Impact of Enforcement on Tax Compliance. *The Journal of Law and Economics 61*(1), 1–35. Publisher: The University of Chicago Press.

Dom, R., A. Custers, S. Davenport, and W. Prichard (2022, February). *Innovations in Tax Compliance: Building Trust, Navigating Politics, and Tailoring Reform.* The World Bank.

Eissa, N. and H. Hoynes (2011, June). Redistribution and tax expenditures: the earned income tax credit. *National Tax Journal 64*(2.2), 689–729. Publisher: The University of Chicago Press.

Eissa, N., H. J. Kleven, and C. T. Kreiner (2008, April). Evaluation of four tax reforms in the United States: Labor supply and welfare effects for single mothers. *Journal of Public Economics 92*(3), 795–816.

Farooq, Z. (2024). Static and Dynamic Impacts of Corporate Tax Audit on Firm Evasion and Compliance Behavior.

Finkelstein, A. and N. Hendren (2020, November). Welfare Analysis Meets Causal Inference. *Journal of Economic Perspectives 34*(4), 146–167.

Gorodnichenko, Y., J. Martinez-Vazquez, and K. Sabirianova Peter (2009, June). Myth and Reality of Flat Tax Reform: Micro Estimates of Tax Evasion Response and Welfare Effects in Russia. *Journal of Political Economy 117*(3), 504–554. Publisher: The University of Chicago Press.

Guyton, J. and R. H. Ii (2021). The Compliance Costs of IRS Post-Filing Processes. *An IRS-TPC Research Conference: Advancing Tax Administration*.

Heckman, J. J., S. H. Moon, R. Pinto, P. A. Savelyev, and A. Yavitz (2010, February). The rate of return to the HighScope Perry Preschool Program. *Journal of Public Economics 94*(1), 114–128.

Hendren, N. (2016). The Policy Elasticity. *Tax Policy and the Economy 30*, 51–89. Publisher: [The University of Chicago Press, The National Bureau of Economic Research].

Hendren, N. and B. Sprung-Keyser (2020). A Unifed Welfare Analysis of Government Policies. *Quarterly Journal of Economics 135*(3), 1209–1318.

Jensen, A. (2022, January). Employment Structure and the Rise of the Modern Tax System. *American Economic Review 112*(1), 213–234.

Kaldor, N. (1963). Taxation for Economic Development. *The Journal of Modern African Studies 1*(1), 7–23. Publisher: Cambridge University Press.

Keen, M. and J. Slemrod (2017, August). Optimal tax administration. *Journal of Public Economics 152*, 133–142.

Kitagawa, T. and A. Tetenov (2018). Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice. *Econometrica 86*(2), 591–616. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13288.

Mayshar, J. (1990, December). On measures of excess burden and their application. *Journal of Public Economics 43*(3), 263–289.

Mayshar, J. (1991). Taxation with Costly Administration. *The Scandinavian Journal of Economics 93*(1), 75–88. Publisher: [Wiley, The Scandinavian Journal of Economics].

Nie, X. and S. Wager (2020, August). Quasi-Oracle Estimation of Heterogeneous Treatment Effects. arXiv:1712.04912 [econ, math, stat].

Olken, B. (2007, April). Monitoring Corruption: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy 115*(2), 200–249. Publisher: The University of Chicago Press.

Reeves, C. R. (Ed.) (1993, April). *Modern heuristic techniques for combinatorial problems.* USA: John Wiley & Sons, Inc.

Robins, J. M., A. Rotnitzky, and L. P. Zhao (1994). Estimation of Regression Coefficients When Some Regressors Are Not Always Observed. *Journal of the American Statistical Association 89*(427), 846–866. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

Saez, E., J. Slemrod, and S. H. Giertz (2012, March). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature 50*(1), 3–50.

Slemrod, J. and S. Yitzhaki (1996). The Costs of Taxation and the Marginal Efficiency Cost of Funds. *Staff Papers (International Monetary Fund) 43*(1), 172–198. Publisher: Palgrave Macmillan Journals.

Stern, N. H. (1976, July). On the specification of models of optimum income taxation. *Journal of Public Economics 6*(1), 123–162.

Sverdrup, E., H. Wu, S. Athey, and S. Wager (2024, April). Qini Curves for Multi-Armed Treatment Rules. arXiv:2306.11979 [stat].

Wager, S. and S. Athey (2018, July). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association 113*(523), 1228–1242. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01621459.2017.1319839.

Waseem, M. (2018, January). Taxes, informality and income shifting: Evidence from a recent Pakistani tax reform. *Journal of Public Economics 157*, 41–77.

Yadlowsky, S., S. Fleming, N. Shah, E. Brunskill, and S. Wager (2023, November). Evaluating Treatment Prioritization Rules via Rank-Weighted Average Treatment Effects. arXiv:2111.07966 [stat].

A Appendix: Additional Figures

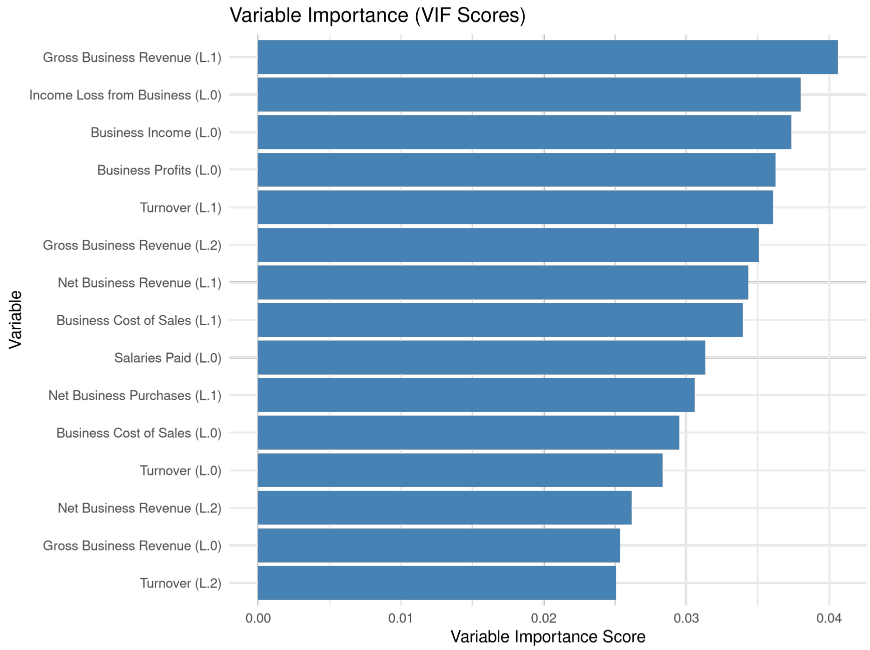

Figure 5: RATE: NPV Revenue

This figure plots the variable importance factors (VIFs) for the top 15 most “important” variables in the causal forest prediction exercise. VIF is computed as the share of trees which split along a given covariate, weighted by the depth at which the split occurred so that splits which occur earlier in trees are weighted more heavily.