On Optimal Fiscal Policies under Heterogeneity

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der KIT-Fakultät für Wirtschaftswissenschaften
des Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

M. Sc. Ezra Sebastian Setio

Tag der mündlichen Prüfung: 16. Dezember 2025

Referent: Prof. Dr. Berthold Wigger

Korreferent: Prof. Dr. Frank Kupferschmidt

Karlsruhe, Dezember 2025

Table of Contents

Chapter 1

Introduction

At which stage in a social planner's welfare analysis or policy design does heterogeneity take place? Where *should* it?

These questions motivate the topics explored in this dissertation and remain as common thread along the research inquiries throughout. The first half of the dissertation, that is, the Chapters 2 and 3, are devoted to the field of welfare analysis. Within the public finance literature, existing methodologies for welfare analyses exist across diverse sectors of the field, with varying emphases and conventions on the preferred metric and methodological approach to evaluating a given policy or reform. On the other hand, proliferation of increasingly rigorous econometric research studying specific policies in at least the past three and a half decades presents the welfare analyst with an abundance of causally estimated findings he can then leverage as components of either the projected benefits a given policy imparts on its beneficiaries or the costs to the government in form of tax-affecting behavioral responses from those beneficiaries. Chapter 2 of this dissertation categorizes these empirical insights into relevant aspects of expected fiscal externalities from a fiscal policy and the corresponding 'out-of-own-pocket' willingness to pay from the impacted individuals using the marginal value of public funds (MVPF) framework, contributing thereby through both empirical expansion and methodological extensions into the previously uncovered taxation and expenditure sectors. Chapter 3 illuminates another aspect of *using* the resulting, unified values of public spending across policy sectors in order to achieve a welfare-relevant comparison.

In the second half of the dissertation, the focus turns more into individual heterogeneity across the population, as represented in for example sample populations in detailed administrative datasets. Recent literature suggests that leveraging knowledge of heterogenous effects (which recent advances in econometric methods supported by statistical and machine learning have increasingly enabled) might not only serve distributional goals *ex-post* (i.e., after a certain public program has been determined or even executed), but might also unlock efficiency gains if incorporated *ex-ante* into the planning of the policy (e.g., as an explicit part of treatment allocation optimization – as in the Policy Learning literature strand of applied econometrics and causal machine learning, represented in Chapters 4 and 5 of this dissertation.

Chapter 2

Welfare Analysis of Fiscal Policies Across Public Sectors

Expanding Methodological and Empirical Aspects of Welfare Analysis Framework Marginal Value of Public Funds

Chapter 2.1 Introduction

The literature on public economics in general and welfare analysis more specifically have both, along with other branches of (applied) economics, benefitted from the dynamic proliferation of empirical methods in the past at least two and a half decades—the period of time in which the methodology of economics research as a whole has seen wide-ranging increase in its empirical techniques, with substantial emphasis on methods that uncover causal variables. Study designs such as instrumental variables (IV), difference-in-differences (DiD, also abbreviated as DD in some parts of the literature), and regression discontinuity represent flagship groups of empirical strategy that explicitly target estimation of *estimable* parameters, i.e., causal variables whose effects on other, dependent variables can be causally estimated in the chosen empirical context. The Nobel Prize in Economics in 2021 to three of the leading economists pioneering this line of approach solidified the widespread of the framework that was sometimes referred to as the “credibility revolution” of econometrics and related fields of applied economics.

Yet despite these advances in the methodology and the subsequent increase in available empirical findings, the speed with which public finance literature can take advantage of these developments and leverage them as crucial elements to answering normative questions could arguably be further accelerated—a point which a recent literature on marginal value of public funds brings forward, while at the same time building upon older traditions of welfare analysis such as marginal cost of public funds and marginal excess burden literatures. Chief among those normative considerations in the context of public economics and especially with regards to fiscal policies are determination of key variables that constitute a *good* economic policy, along which the optimality conditions under which such a policy can be classified as being *best* among available alternatives. Broadly speaking, it is to the expanding of the current research state-of-the-art and to the better understanding of those questions that Chapter 2 endeavors.

Chapter 2 and Chapter 3 utilize as welfare analysis framework the Marginal Value of Public Funds (MVPF). While essentially a continuation of the welfare analysis tradition that includes, among others, the marginal cost of public funds and marginal excess burden—both of which will be discussed in relation to the MVPF in the next chapter, the MVPF method has since its introduction in Hendren & Sprung-Keyser (2020), seen expanded use in the public economics literature. Of didactic and research value is at the same time the Policy Impacts library[1] which acts as a collaborative platform that complements traditional publication channels in gathering calculated welfare impacts of a wide array of public policies from different countries.

This Chapter reports implementations of the MVPF framework in two sectors of public expenditure policies: active labor market programs (ALMP, with job training as its most common representative) and social insurance policies. The marginal value of public funds can be used to compare the value of same amount of government budget directly against the expected value of a direct transfer of equivalent amount to either the same beneficiary group (more on this particular comparative function in Chapter 3) or to other beneficiary group (thus necessitating the explicit stating of the planner's assumed social welfare weights/preferences for redistribution). The remainder of Chapter 2 is devoted to illuminating several taxation forms that have previously eluded MVPF analysis in the literature.

[1] The Policy Impacts Team at Massachusetts Institute of Technology maintains a continuously expanding collection of MVPF results from across the world, accessible at policyimpacts.org.

Chapter 2.2 Related Literature

The marginal value of public funds framework is most closely related to the following welfare evaluation concepts: 1) the marginal efficiency cost of funds as defined in Slemrod and Yitzhaki (1996) and Slemrod and Yitzhaki (2001); 2) the marginal cost of public funds in the definition of Kleven and Kreiner (2006). as well as Mayshar (1990)'s definition of marginal excess burden[2]. In particular, (Slemrod & Yitzhaki, 1996, 2001) give theoretical microfoundations to quantifying the welfare impacts of *increasing* tax rates, which they call the *negative* marginal efficiency cost of funds. As will later be described more thoroughly, this line of argumentation is closely related to the implementation of MVPF for top income tax rates in previous literature and the novel extensions of the MVPF formula into three types of capital taxation given in Subchapter 2.3.3.

In regards the use of *causal estimates* from empirical studies using *causal identifying strategies* to determine the *policy elasticity* as the direct causal measure of the behavioral impact *without needing* to decompose into (structurally founded, not-directly estimable functional form assumptions needed) *substitution effects* and *income effects*, Hendren (2016) contains the definition of policy elasticity and the corresponding theoretical foundations for its direct usage as measures of willingness to pay and fiscal externalities in the MVPF framework. For overview on which empirical studies fulfill the causality criteria needed to be incorporated as elements of the MVPF being calculated, see, e.g., the widely-used handbook by Angrist and Pischke (2009), as well as the first author's and his fellow Nobel laureates' acceptance lectures at the 2021 Nobel Prize. Also closely related (and frequently assumed in different variations across empirical studies of each policy sector) is the literature on *sufficient statistics* (Chetty, 2009; Kleven, 2021)

[2] Subchapter 4.2 goes deeper into how the MVPF differs, on the other hand, from older definitions of marginal cost of public funds and marginal excess burden elsewhere in the public finance literature.

Also related to the MVPF literature is the rather diverging view offered by the methodological literature focusing on the use of structural modellings. For example, García and Heckman (2022) proposes Net Social Benefits (NSB) as a counterpart to the MVPF. In this regards, I note that the discourse appears to continue in the literature Hendren & Sprung-Keyser (2022), yet does not constitute the core topic in the scope of this dissertation.

Finally, the use of MVPF results in the unified framework to compare policies across government sectors represents an aspect which, to my view, has enjoyed a less than proportionate amount of attention in the literature so far and thus merits a more intensive treatment in this dissertation. My exercises in shedding light into these interpersonal and inter-beneficiary-group comparisons using existing library of MVPF results, as well as the related literature thereto, can be found thus in Chapter 3.

Chapter 2.3 Implementations and Results

The following subchapters delve into each sector and its corresponding MVPF calculation. This partially includes, where applicable, the necessary methodological extension for the type of fiscal policy being studied. The Subchapter 2.3.1 begins with public expenditures for the labor market at the participation margin, i.e., job training and related active labor market interventions.

Chapter 2.3.1 MVPF of Active Labor Market Policies (ALMP)

Caliendo, Hujer and Thomsen (2008) showed that while prior studies found negative average treatment effect for German job creation schemes (JCS), disaggregating enabled them to present a clearer picture of the situation. In other words, they discovered no negative effects for most segments yet singularly positive effect for the long-term unemployed.

The literature on comparative efficiency of active labor market programs is extensive, with several surveys providing excellent summary (e.g., Card et al., 2010, 2018; Kluve, 2010; Vooren et al., 2019). Most recently, McCall et al. (2016) delivers a comprehensive survey on particularly job training policies and reforms in the four major advanced economies of United States, United Kingdom, France and Germany, while Le Barbanchon et al. (2024) went further by juxtaposing the the generated impacts of these reforms with those of unemployment insurance (UI) policy changes targeting similar beneficiary groups. In particular, German active labor market programs tend to generate positive effects for the longer run, albeit taking a relatively long period in the beginning to achieve break-even point[3] (Lechner et al., 2011). The gain manifests itself in either employment chances, earnings increase, or both.

[3] This is known in the labor economics literature as lock-in effect (e.g., Lechner et al., 2011).

In the case of publicly administered job training, benefit to participating individuals is measured by income premium. Empirical estimates of individual earnings gain caused by participating in the job training program constitutes thus the main part of willingness to pay. Subtracted from this value is then ideally any reduction in cash transfers (such as unemployment allowance) that these participants used to receive but to which they are no longer entitled, due to their current participation in job training program. This approach is based on the envelope theorem which assumes that the change in budget constraint is a pure result of the policy, i.e., instead of being a joint result of the policy impact and responses in labor supply.[4] For its part, net cost to the government per program recipient is generally provided in the empirical studies and is relatively straightforward.

In light of this recent development in the literature chiefly positioned in the overlap between labor economics and public finances, there are significant advantages of incorporating the empirical findings (often reported hitherto through varying measures of efficacy: the variants of Baily-Chetty formula adapted from the public insurance literature (for example in McCall et al., 2016 and the empirical papers on vocational and adult professional education trainings referenced therein); specific quantifications of the participation constraint and thus effects at the extensive margin (thereby focusing on the *marginal workers*, namely the precise segment of the labor market whose decision on participation into/out of the labor market can with highest probability be affected by the reform under study); or, more conventionally, using net present values (NPV) of the projected changes in income and other utility-affecting levers observable to some degree across the beneficiary group (with contributions from both sufficient statistics estimated using instrumental variable (IV) designs *and* from prominent representatives of structural evaluation methods).

This convincing proliferation of empirical estimation results notwithstanding, the benefits to be recuperated from integrating these findings into a unifying welfare

[4] For theoretical foundations of envelope theorem in measuring willingness-to-pay of recipients of fiscal and regulatory public policies, consult for example (Finkelstein & Hendren, 2020).

framework – i.e., the welfare-analytical arguments brought forward in preceding Subchapters – remain a 'low-hanging-fruit' for research following up on precisely this intersection between literatures of public and labor economics. As the advantages of the marginal value of public funds (MVPF) framework as one methodological tool toward this end has been elucidated before, the remaining of Subchapter 2.3.1 lends itself to illustrating the usage of the method by means of two empirical applications for active labor market programs (ALMPs) in Germany across the past five decades[5].

Germany's public job trainings 1975–1997

Germany's Research Data Centre (*Forschungsdatenzentrum*, FDZ) Federal Institute for Labor Market and Employment Research (*Institut für Arbeitsmarkt- und Berufsforschung*, IAB) maintains several comprehensive dataset available on research request[6] pertaining to, among others, evaluating publicly administered training programs. Lechner et al. (2011) utilizes a version of these covering the intervention periods 1975–1997, highlighting retraining programs constituting of on average 21 months of training.

Short trainings refer to job training programs finishing after less than or at most 6 months and generally display shorter lock-in effect as well as lower monthly cost. Participating individuals achieve positive net gain after on average 1 year and accumulates an average of €42,000 earnings surplus 8 years in the future, whereas the cost to the government is estimated at €4,439 per month (including no additional cost for accumulated unemployment). Taken together, these figures assemble an MVPF of $\frac{\$42{,}000}{\$4{,}439} = 9.46$ for government-supported short training programs.

[5] A version of these calculations were reported in (Setio, 2021). This Subchapter expands those results by incorporating novel findings in the literature on *benefit projections* as recently summarized by (Le Barbanchon et al., 2024) and therefore, naturally, the corresponding changes in *fiscal externalities* calculation.

[6] The oldest-available of which, the SIAB dataset, can be for example seen in Chapters 4 and 5 of this dissertation—there in utilization concering statisical decision theoretical method called Policy Learning.

Publicly financed job trainings are also the subject of Bernhard's (2016) investigation, where he documented similarly short lock-in effect for short-term training programs with participants experiencing on average €274 reduction to their monthly income for the first 6 months but benefiting from €286 more income per month for the rest of the 104-month-long observation period. I use Lechner et al.'s (2011) cost estimate and align the observation period with theirs (8 years) for more accurate comparability, arriving at

$$\text{MVPF} = \frac{-€274 \times 6 \ + \ €286 \times 90}{€4{,}439} = 5.428$$

For long term trainings, Bernhard (2016) reported an average €340 monthly loss of income for the first 2 years, after which an average participant benefit from €416 monthly earnings premium. Note that while Bernhard (2016) assigned the label *long-term* for all observed trainings lasting more than 6 months, Lechner et al. (2011) distinguished them further into *retraining* and *long* trainings. Bernhard (2016) explicitly related his findings for long-term trainings with Lechner et al.'s (2011) accounts on retraining programs. In line with more recent findings (cf. McCall et al., 2016), the average cost of *retraining* can thus be estimated as the €20,983 average per participant cost for my MVPF calculation:

$$\text{MVPF} = \frac{-€340 \times 24 \ + \ €416 \times 72}{€20{,}983} = 1.039$$

The One-Euro-Job program

The *Ein-Euro-Job* program is a job assistance program instituted in Germany between 2005–2007. Targeting individuals at the participation margin of the labor force, the program administers complementary job for recipients of unemployment benefit (*Arbeitslosengeld II)*. Using administrative data from the Institute for Employment Research (IAB), Harrer and Stockinger (2019) evaluated the effects of One Euro Job program after its reform in 2012.

The One Euro Job program had an average monthly cost per participant of €459, of which €124 is cash transfer to the participant in the form of lump-sum expense allowance (*Aufwandsentschädigung*). Because participation in One Euro Job program does not automatically impart ineligibility for further or re-subscription to unemployment benefit, there is no cost reduction involved due to spared unemployment benefit expenditure (as is the case in some job training programs used in other MVPF calculations. Furthermore, since the transferred lump-sum allowance is in average (€1,488 yearly) far below the annual tax-exempt amount (*Grundfreibetrag*, in 2013 equals €8,130), the net cost of One Euro Job program does not need to incorporate any change in tax contribution.[7] The average participation duration was 4 and a half months, while the maximum attainable duration was 12 months.

My MVPF calculation consequently reflects these two possible scenarios. However, in line with most recent findings summarized in Le Barbanchon, Schmieder, and Weber (2024), calculating the costs toward the longest attainable/extendable duration has been recommended, in order to better highlight the interplay between job training programs and unemployment insurance—whose very recipients are increasingly confirmed by most recent studies to display consistent overlap. As a result, I highlight here the second scenario.

On the other side of the scale, willingness to pay for One Euro Job program is derived from the received work compensation (€124 monthly on average) and change in earnings. Harrer and Stockinger (2019) reported negative effect on participants' earnings for 3 years after the program's start, ranging between –€330 to –€110 for residents of former West Germany states and –€220 for former East Germany states.

[7] To be more precise, €1,488 represents average annual earnings during program participation which lasts at most one year, but it serves as ceiling for second and third-year income approximation given the reported persistent negative income trend among this cohort.

Putting these elements together, the marginal value of public funds for the One Euro Job program can be estimated as follows.

$$\text{MVPF}_{\text{ALMP}} = \frac{\text{periodical allowance}\left[\frac{€}{t}\right] \times \text{duration [t]} - \text{annual earnings loss [€]}}{\text{program cost}\left[\frac{€}{t}\right] \times \text{duration [t]}}$$

Correspondingly, calculating for program participation with maximum duration (one year) estimated by Harrer and Stockinger (2019)[8] gives us:

$$\text{MVPF} = \frac{€124 \times 12 - €220}{€459 \times 12} = 0.23$$

[8] Note that Harrer and Stockinger (2019) did not specify substitute revenue source (if any) after program exit, meaning the only changed MVPF component if one were to extend evaluation period unto their 3-year timeframe would be multiplicated yearly earnings loss.

Chapter 2.3.2 MVPF of Social Insurance Policies

In this subchapter, I report calculations of the marginal value of public funds for each of four sectors of public social insurance, beginning with health insurance, followed by retirement/old-age, unemployment and disability insurance[9]. For these sectors of public expenditure, the state of the art of the public finance literature features two particular strands of development that have recently been widely accepted to be common practice (or at least commonly integrated as a benchmark analysis/mechanism consideration part of the studies): the incorporation of ex-ante measures of willingness to pay and the value of offering choice when mandate is the status quo. Feasibility to incorporate these advancements into the standard MVPF formula is illustrated in the following empirical examples, while more discussion on their current limitations can be found in Chapter 2.4.

The next sections deliver empirical extensions of the two novel methodological innovations in welfare impact evaluation of the literature on social insurances. To do so, I first revisit my own findings for the marginal value of public funds of several historical policies and/or reforms of publicly administered insurances (Setio, 2021)—the calculations of which were previously mainly based on variants of the what is commonly known in the literature of public insurances of the recent decades as Baily-Chetty formula (for my derivations of and the corresponding quantitative results using the aforementioned non-modified MVPF formula).

Therefore, before delving into the MVPF extension with measurement of ex-ante willingness to pay, I present in the following a recap of the key elements for the original MVPF calculations (cf. Setio, 2021).

[9] The latter three categories of state-administered social insurances can display varying degree of overlap across different countries or even regions within them, and/or have been analyzed combinatively in previous studies (e.g., Landais et al. 2021; Spinnewijn 2020)

MVPF for unemployment insurance reforms

The ratio between behavioral and mechanical costs (henceforth BC, MC) is hitherto commonly used to evaluate the welfare effects of an unemployment insurance policy (e.g., Schmieder & von Wachter, 2017; Ye, 2018). The rationale for adopting BC/MC ratio in unemployment insurance context can be compared to the motives behind the success of elasticity of taxable income (ETI) as a measure of fiscal externalities due to a given tax change. ETI captures behavioral responses to a tax reduction/increase in a more comprehensive way than the traditional labor supply elasticity. For its part, BC/MC ratio enhances the ability of conventional labor supply elasticity in capturing behavioral responses to rising/falling unemployment insurance benefits (see also Saez et al., 2012; Schmieder & von Wachter, 2017).

Recall the earlier established notion that for a marginal policy change, willingness to pay is defined as equivalent to mechanical cost, assuming the envelope theorem holds (following this assumption, any change to the individual's budget constraint is attributable to the policy change, and not to changes in labor effort). In other words, the MVPF can be equivalently defined as[10]

$$\text{MVPF} = \frac{\text{Mechanical cost}}{\text{Mechanical cost} + \text{Behavioral cost}}$$

Normalizing for mechanical cost gives an equivalent definition as follows.

$$\text{MVPF} = \frac{1}{1 + \frac{\text{Behavioral cost}}{\text{Mechanical cost}}}$$

Analyzing unemployment insurance programs in 72 countries, Schmieder and von Wachter (2017) found that there is no consensus on whether benefit level is more effective than extending program duration: for the average of countries within the Organization for Economic Cooperation and Development (OECD), a benefit level increase proves more effective; for the bottommost quantile, an extension in

[10] Relatedly, Ye (2018) expresses marginal cost of efficiency funds as $1 - \frac{BC}{BC+MC}$.

program duration is marginally more effective. With a BC/MC ratio of 0.35 (Schmieder and von Wachterm 2017) for the average across OECD countries and using the last reformulation of MVPF, the benefit level increase yields an MVPF of $\frac{1}{1+0.35} = 0.7407$. For the duration extension policy with a BC/MC ratio of 0.58 (Schmieder and von Wachter 2017), the corresponding marginal value of public funds equals $\frac{1}{1+0.58} = 0.6329$.

MVPF for health insurance reforms

Using 2011–2014 LIFE-Adult-Study dataset from Leipzig Research Centre for Civilization Diseases, Hajek et al. (2020) reports a monthly average willingness to pay of €240 for health insurance, which represents approximately 14% of average net monthly household income in Germany. German senior citizens studied up to 2010 display a slightly higher monthly willingness to pay for health insurance at €261, which in turn constitutes 18% of their €1,433 average disposable income (Bock et al., 2016).

Federal Statistical Office's official health system databank Federal Health Monitoring (Gesundheitsberichterstattung des Bundes, 2021) documented an annual average government expenditure per public-health-insured individual of €2,090 in 2011 and €2,355 in 2014.

I combine each of the latter two with Hajek et al.'s (2020) willingness-to-pay estimate to form

$$\text{MVPF}_{\text{Average 2011–2014}} = \frac{€240 \times 12}{€2223} = 1.296\ .$$

For seniors, their higher average willingness to pay is reflected in:

$$\text{MVPF}_{\text{Seniors}} = \frac{€261 \times 12}{€2{,}057} = 1.523$$

MVPF for retirement insurance reforms

In Germany, disability insurance is normally discussed in combination with retirement insurance. A comprehensive review of the historical development of German retirement and disability insurance is given in Boersch-Supan and Juerges (2011) (see also Boersch-Supan & Schnabel, 1998). In the existing literature, studies that report empirical findings on welfare effects of German retirement and disability insurance policy changes have employed different strategies to estimate individual welfare gain or loss.

Germany started in 2014 granting exceptions to the aforementioned early retirement income deduction, namely, individuals who have contributed to the pension scheme for at least 45 years were then given the option to enter retirement at age 63, which is four years early than the normal threshold. Using administrative dataset *Versichertenrentenzugang* between 2013–2017, Krolage (2020) reported total yearly increase in transferred pension benefits due to new retirements-at-63 as well as accompanying opportunity costs to the government in the form of foregone social contributions and tax payments. While Krolage's (2020) estimate for pension insurance expenditures represents MVPF's benefit to the beneficiaries, her estimate for total costs neatly translates to MVPF's net cost because it includes fiscal externalities in form of lost social contributions and tax revenue. Note that while technically MVPF refers to individual willingness to pay and net cost, in this case it would mean dividing both the numerator and denominator each (because Krolage's (2020) both estimates already refer to the exact same group of beneficiaries) with the number of at-63 retirees (e.g., 225,290 persons in 2016) and thus it is not necessary to reformulate the aggregate values. For example, the MVPF of exempting long-term-contributors from early retirement deductions was in 2014:

$$\mathrm{MVPF}_{2014} = \frac{€0.82\text{ billion} \times \frac{1}{\text{Number of retirees 'at 63'}}}{€1.51\text{ billion} \times \frac{1}{\text{Number of retirees 'at 63'}}} = 0.54$$

MVPF extended with Ex-Ante Willingness to Pay

Following the current state of the literature strand (see Einav et al., 2010 for review; Finkelstein et al., 2019 for an empirical application of MVPF with ex-ante willingness to pay; and Hendren et al., 2021 for theoretical underpinnings of the ex-ante willingness to pay approach), the availability of a contextualized measure of *survey-elicited* willingness to pay for insurance of *the insured* individuals, D(s). While the calculation of such measure for other countries and sectors of social insurances remains an undoubtedly urgent and highly relevant avenue for further research[11], here I proceed for purpose of methodological development with updating several MVPFs for health insurance reforms (thus of which the beneficiaries can be expected to display similarity– to a degree that reflects the availability in the current literature on public health insurance. The modified MVPF with ex-ante measure of willingness to pay can be written as (Hendren, 2021):

$$MVPF_{\text{Ex-Ante WTP}} = \frac{1+(1-s)*\gamma(D(s)-E[D(s')] \mid s' \geq s]}{1-\frac{D(s)-C(s)}{s(D'(s))}}$$

, which represents a slightly simplified equation form of the one derived in (Hendren, 2021) specifically for incorporating the *value* the recipient places 'from behind the veil of ignorance' – that is, the premium they are revealed to actually be willing to pay for the insurance *before* further information on their true risk becomes available (i.e., *ex-ante welfare* or *ex-ante willingness to pay*). Here, $s \in [0,1]$ denotes the proportion of the population covered by the state insurance; $C(s)$

[11] On the one hand, one such estimation arguably necessitates a rather comprehensively planned and high-continuation time-panel study, potentially involving a new survey eliciting individuals's private valuation (e.g., Finkelstein et al., 2019). On the other hand, once calculated, the result arguably offers multifold research benefits, as it could then be more straightforwardly combined with preceding studies for countries for which results on γ (risk aversion) and *C(s)* (average cost) already exist separately.

denotes the costs imposed on the state as insurer at the s proportion of the insured and $D(s)$ their corresponding willingness to pay[12].

In the numerator of this MVPF calculation, the additional term $(1-s) * \gamma(D(s) - E[D(s')]|s' \geq s]$ captures the ex-ante willingness to pay (the additional risk-coverage premium the marginal beneficiary—those whose induced-entry into the insurance due to the policy would increase the insurance size from the status-quo s to a new level $s' \geq s$—would value had her expected utility been measured prior to the revelation of her true risk, or in other words) of the as-yet uninsured part of the population, i.e., $(1-s)$. $D(s) - E[D(s')]$ thus gives the average difference of willingness to pay between the already-insured and the marginal beneficiary. This difference is then multiplied by a risk-coefficient factor γ, which I follow Hendren (2021) in setting to 0.0005 in accordance to the common estimate in the health insurance literature. Consistent with the average WTP and cost for the insured, I here take Finkelstein et al.'s (2019) findings for $D'(s)$ as nearest approximation. Future studies could undoubtedly benefit from more contextualized estimates of these three variables, provided they are armed with data on random price variation for each setting of the respective social insurance reform.

Reformulating this modified MVPF formula mathematically leads to an equivalent formulation, now re-stated as the product of the standard MVPF[13] (i.e., not yet extended with ex-ante valuation—in other words, the MVPF on health insurance policies measuring *observed/ex-post willingness to pay*), as follows:

$$MVPF\ \textit{Ex-Ante WTP} = MVPF_{Standard} * (1 + (1-s) * \gamma(D(s) - E[D(s')|s' \geq s])$$

[12] $D(s)$ and $C(s)$ are estimated through the Einav-Finkelstein framework (see, for review, Einav & Finkelstein, 2023), which represents currently the benchmark approach for estimating in the social insurances literature. Due to unavailable data of random price variation for each of the five selected reforms, I use the Finkelstein et al.'s (2019) calculations as nearest approximations.

[13] As in Chapter 3 and in Setio (2021), MVPF for social insurance settings is calculated through a reformulation of MVPF $= \frac{1}{1 + \frac{\text{Behavioral cost}}{\text{Mechanical cost}}}$.

To give an example, combining estimates across OECD countries from Schmieder and von Wachter (2017) of a benefit-level-increase reform in unemployment insurance *and* the additional premium from ex-ante willingness to pay at an assumed insurance level of 30%[14],

$$\begin{aligned} MVPF\ \textit{Ex-Ante WTP} &= 0.74 * (1 + (1 - 0.3) \times (5 \times 10^{-4}) \times (1978 - 853]) \\ &= 1.03 \quad . \end{aligned}$$

Similar procedure was applied for health insurance reforms of Germany as previously reported in (Setio, 2021), which was however silent on the possibility of incorporating the ex-ante valuations. Since the “Join the Healthy Boat” is a very narrowly targeted, sample-means-tested study in one federal state in Germany (Baden-Württemberg), following the guidelines established earlier I chose to focus the calculation of the extended MVPF on the general adult health insurance results. The main results are reported in the Table 2.1.

Policy sector	Program type	MVPF Standard	**MVPF Ex-Ante WTP**
Unemployment insurance, OECD mean	Benefit level increase	0.74	**1.03**
	Duration extension	0.63	**0.88**
Health insurance, Germany	Health insurance adult	1.30	**1.81**
	HI Adult, Seniors only	1.52	**2.12**
Retirement insurance, Germany	Introduction of early retirement threshold	0,54	**0.75**

Table 2.1 Non-modified MVPF results and new MVPF results supplemented with ex-ante measurement of willingness to pay.

[14] Following the baseline adopted in Hendren (2021).

Chapter 2.3.3 MVPF of Capital Taxations

Turning to the public revenues side[15], this subchapter derives novel MVPF formulas for three forms of capital taxation particularly present in the current public finance discourse (for recent reviews consult (Bastani and Waldenström 2020; Jakobsen et al. 2020; and Piketty, Saez, and Zucman 2023). Institutionally, the three forms of capital taxation have experienced varying trajectory in the past several decades, with wealth taxes and estate taxes having been abolished and to some extent re-introduced in several major economies, whereas the capital income tax underwent dynamic considerations highlighting the interaction with public discourse on the taxation of (un-)realized capital gains.

MVPF of Inheritance and/or Gift/In-Vivo Tax Reforms

In the following, I derive a marginal value of public funds (MVPF) formula for reforms of inheritance taxes. This effort contributes to the literature of welfare analysis a *methodological novelty*, because the current state of the literature exhibits a research gap of welfare impact quantification[16] for inheritance tax reforms. There exists a moderate body of literature laying theoretical grounds for calculating *optimal* (i.e., normative) inheritance tax rates (García-Miralles, 2020; Piketty & Saez, 2013), yet this literature strand does *not* consider welfare impact calculation of *historical* (i.e., positive/evaluative) inheritance tax reforms.

On the other hand, a small but growing strand of literature investigates *behavioral responses* in specific contexts of inheritance tax changes. Their findings are valuable, but rely unfortunately so far on estimation methods whose

[15] The MVPF framework is suitable both for tax rate cuts (i.e., tax cuts as a form of direct transfer instead of in-kind policies such as the public expenditure programs discussed earlier) *and* for tax rate hikes (related to the earlier literature's *negative* Marginal Efficiency Cost of Funds, see Slemrod and Yitzhaki (1996, 2001).

[16] This research gap pertains not only to the MVPF but more generally to the welfare economic literature, where neither older welfare measures such as marginal efficiency cost of funds (MECF) nor frameworks based on structural variables such as Net Cost Benefit (García and Heckman 2022) offer quantification method for measuring the impact of an inheritance tax reform.

extrapolatability into general external validity remain less proven compared to more established methods such as DiD or IV—studies such as (Glogowsky, 2021 in the context of Germany's coupled inheritance and gift (in-vivo) taxes and Goupille-Lebret and Infante, 2018 for France's estate tax rely solely on "bunching" estimation method).

I derive the formula for marginal value of public funds of a change in a inheritance/estate tax rate by going back to the theoretical foundations underpinning the simplified version of MVPF formula now commonly utilized for income tax rate. This procedure also reflects how the MVPF formula for top income tax rate was originally established. I first approach the *elasticity of taxable bequest* and denote how the elasticity of taxable bequest constitutes *behavioral response* of the inheritance taxpayer (depending on the specification of bequest motive, this could be the bequest leavers as well as the bequest receivers or weighted combination thereof) in response to a marginal change in the *net-of-tax rate* (i.e., changes in negative proportion to the change in tax rate) derived in a manner of sufficient statistic (estimable from empirical variables, independent of functional form assumptions):

$$\varepsilon_B = \frac{dB}{B} \cdot \frac{1 - \tau_B}{d(1 - \tau_B)}$$

I proceed by using this elasticity of taxable bequest to calculate the effect of $d\tau_B$, the change in inheritance tax rate, on total tax revenue expected by the state planner, dR. Consistent with the larger public finance literature (for comprehensive reviews see Diamond and Saez (2011) and Saez, Slemrod, and Giertz (2012), the latter term consists of, on the one hand, the *mechanical* revenue effect dM (i.e., the marginal amount of foregone revenue a government expects in the case of a marginal tax cut) and the *behavioral* revenue effect dB:

$$dR = dM + dB$$

Adopting the standard definition of Pareto parameter for income distribution in the public taxation literature (e.g., Atkinson et al., 2011; Diamond and Saez, 2011), $\alpha = \frac{E[y_i \mid y_i \geq \bar{y}]}{E[y_i - \bar{y} \mid y_i \geq \bar{y}]}$ (with $\bar{y}$ representing the income threshold over which the income tax rate under study applies and y_i the income of the marginal individual for whom the MVPF is calculated), I derive Pareto parameter distribution for bequest level b_i at bequest threshold $\bar{b}$ as $\alpha_B = \frac{E[b_i \mid b_i \geq \bar{b}]}{E[b_i - \bar{b} \mid b \geq b]}$.

Re-formulating further,

$$
\begin{aligned}
dR &= dM + dB \\
&= (B_i - \bar{B}) \, . \, d\tau_B \; - \; \tau_B \, . \varepsilon_B \, . B_i \, . \, \frac{d\tau_B}{1-\tau_B} \\
&= (B_i - \bar{B}) \, . d\tau_B \, . \; \left[1 \; - \; \varepsilon_B \, . \, \frac{B_i}{B_i - \bar{B}} . \, \frac{\tau_B}{1-\tau_B}\right] \\
&= (B_i - \bar{B}) \, . d\tau_B \, . \; \left[1 \; - \; \varepsilon_B \, . \, \alpha_B . \frac{\tau_B}{1-\tau_B}\right] \\
&= dM \, . \left[1 \; - \; \varepsilon_B \, . \, \alpha_B . \frac{\tau_B}{1-\tau_B}\right],
\end{aligned}
$$

where $dM = (B_i - \bar{B}) \, . \, d\tau_B$

and $dB = \tau_B \, . \, dB_i = -\tau_B \, . \varepsilon_B \, . B_i \, . \, \frac{d\tau_B}{1-\tau_B}$

Finally, the marginal value of public funds for inheritance tax reforms can be calculated as follows:

$$
MVPF = \frac{dM}{dR} = \frac{dM}{dM \, . \left[1 \; - \; \varepsilon_B \, . \, \alpha_B . \frac{\tau_B}{1-\tau_B}\right]}
$$

$$
= \frac{1}{1 \; - \; \varepsilon_B \, . \, \alpha_B . \frac{\tau_B}{1-\tau_B}}
$$

To give an example, I apply the formula to Germany's reform of its inheritance tax schedule in 2009. The most relevant estimates available for the necessary components to calculate the formula are found in the studies cited in the following table. The results of the MVPF calculation are also reported at the rightmost column.

As a benchmark, Kopczuk (2013)'s unifying estimate of elasticities related to estate taxation in the US ranges between [0,1 ; 0,2] and is used in the last row to calculate the average impact of a change in US estate tax (with the distributional parameters adjusted to the US setting in the formula).

Elasticity type	Estimated elasticity	Source of estimates of elasticity	Country	Pareto paramete	MVPF
Elasticity of inheritance taxes	0.012 *	Glogowsky (2021)	Germany	1.67	**1,006**
Elasticity of inheritance taxes	0.005 **	Glogowsky (2021)	Germany	1.67	**1,003**
Elasticity of *inter vivo (gift)* taxes	0.012	Glogowsky (2021)	Germany	1.67	**1,006**
Elasticity of inheritance taxes	[0.1 ; 0.2]	Kopczuk (2013)	US	2.29	**[1.078 ; 1.170]**

Table 2.2 *MVPF for reforms on inheritance, estate or gift (in-vivo) tax*

MVPF of Capital Income Tax Reforms

Akin to the derivation for inheritance tax reforms, the following describe the derivation of an MVPF formula for capital income tax reforms under corresponding assumptions (for related theoretical foundations, see Saez and Stantcheva, 2018). To avoid expository redundancy, similar steps are shortened, without loss of generality.

$$MVPF_{Capital\ income\ tax,\ \ 1\%\ rate\ change} = \frac{1}{1 - \frac{\tau_K}{1 - \tau_K} \cdot \boldsymbol{\alpha_K} \cdot \ \varepsilon_K}$$

For capital income tax changes, several empirical results exist in current literature with regards to the *elasticity of taxable capital gains* with respect to the marginal change in the net-of-tax-rate of capital gains tax. For example, Lefebvre, Lehmann, and Sicsic (2025) reports an elasticity result of $\varepsilon_K = 0.7$ for French capital income tax reform, while Agersnap and Zidar (2021) found a remarkably different magnitude of elasticity for United States with an $\varepsilon_K = 1.87$. Excellent public administrative data availability and access for research have for several decades enabled Scandinavian countries to very well-represented in the empirical literature. Fort he elasticity of taxable capital income gains, two distince results are available for Denmark: Kleven and Schultz (2014) arrived at an estimate of $\varepsilon_K = 0.278$ based on 1987 Danish capital income tax reform, while Jakobsen Jakobsen, Kleven, and Zucman (2020) analyzed the 1989 Danish wealth tax reform and obtained an estimate of $\varepsilon_K = 0.486$.

As can already be seen from this selection of results, the empirical literature is still growing and as of now consists of rather wide-ranging interval of elasticity estimates (especially compare this to rather compact estimate for estate taxation given by Kopczuk (2013), between 0.1–0.2). Nevertheless, preliminary implementation of the MVPF formula for capital income gains yield the results of

$MVPF_{Capital\ income\ tax}$ = 2.03 for France, $MVPF_{Capital\ income\ tax}$ = 2.8 for the US, and an average of $MVPF_{Capital\ income\ tax}$ = 1.39 for Denmark[17]

[17] Unlike in the case of income tax and inheirtance tax, calculations for capital income taxes still suffer from unavailability of exact Pareto parameter of capital income distribution for Denmark and France. as pinpointed in the equation above. These results thus will need to be revisited and are of comparatively more preliminary nature.

2.4 Concluding Remarks

This chapter gave account to my methodological contributions to the implementation of the comparative welfare analysis framework of marginal value of public funds in the sectors of inheritance taxation, capital income taxation and the wealth taxation. Quantitative results were also given to the MVPF treatment on German active labor market policies previously evaluated in Setio (2021), where the results are updated by means of additional findings that have since been available in the labor economics and social insurance literature. Lastly, newly-developed insights into calculating willingness to pay of social insurance recipients beyond the traditionally observed market surplus were incorporated into modified MVPF formulas for health and other social insurances.

As elucidated in Subchapter 2.3.2 on the value of public funds for policies and reforms undertaken with regards to public insurances, two branches of the literature represent the state-of-the-art of the methodological and theoretical development in the welfare analysis: along with extension of the market surplus as hitherto convention on quantifying willingness to pay (i.e., the benefits expected by either mandated or self-selected individuals) with *ex-ante willingness to pay* for the risk (arguably the very utility of being insured in the first place, absent extensive additional private information on illness or unemployment risks), the second growing literature strand concerns the (re-)introduction of *choice* (of *supplemental* coverage and/or *differentiated level* of insurance).

In addition to efficiency benefits, public health system offers redistributive advantages. First, universal health care financed by progressive income tax will redistribute income from the rich to the poor. Necessary condition for this relation to hold is sufficiently similar individual preference/'taste' for health – given prescriptions and treatments can only be made by doctors and there exist standard medical practices, this condition is most likely to hold. Secondly, public health care redistributes from the healthy to the sick (or from the young to the aged). Individuals with higher morbidity would have to pay more expensive insurance in the private market. Thirdly, because the rich have the option to opt-out, i.e., subscribe to private insurance (where they can get privileged service), they leave more resources in the public health care to be distributed to the poor (the rich have had to pay the taxes regardless).

The introduction of such choice beyond the one-size-fit-all mandate of social insurance theoretically allows for several different possible dimensions along which the choice could be differentiated. For instance, facilitating an individual at the margin to opt out (or opt in, in less frequent cases where mandates had been rescinded) in a timeframe that becomes increasingly narrower to the expected realization of necessitates an explicit trade-off to be weighed in against the potential increment of adverse selection—an idea that dates back to an early theoretical contribution by Hirshleifer (1971).

With that being said, the fact that involving measures of ex-ante willingness to pay reveals individuals' higher true preferences 'from behind the veil of ignorance' (prior to knowledge of risk) implicitly lends the modification of MVPF with ex-ante willingness to pay, as the calculations in Subchapter 2.3.2, toward being one of those dimensions along which an additional "choice" (beyond the universal mandate) can be identified and estimated. Seen through this conceptual lens, in other words, the context of MVPF extended with ex-ante-WTP represents an area of implementation where the two aforementioned methodological extensions converge into an overlap[18]

Future research would benefit from extending further the empirical collection of MVPF in these major sectors of public expenditures and revenues. Additionally, several related branches would in my view serve particularly well as another methodological extension: the field of property taxation and the field of infrastructure, to name but two. Abundance of cost-benefit analysis variants

[18] Implications derived from this line of argumentation can be generalized into other contexts of social insurances. Moreover, the additional willingness to pay elicited from the individual by way of introducing ex-ante measures also represents incorporation of redistributive effects oft he insurance. Landais et al. (2021) offers in-depth intuition and explanation into both these insights.

influenced by methods of industrial engineering and business economics render the latter a promising avenue to translate *comprehensively* into MVPF framework for better comparisons with similar amount of public spending elsewhere, while the literature of property taxation has seen increasing contribution[19] beyond the literature of public finances, namely from literatures of urban economics, labor economics (relating chiefly to wage-mobility cross-price effects).

Unifying framework such as MVPF also has the potential benefits of suggesting directions where future research avenues can be particularly fruitful. For example, as we saw in the subchapter for tax reforms, more future studies identifying behavioral responses to inheritance tax reforms using empirical methods other than bunching for the same region of tax schedule and the same reform can offer valuable comparison. These, in turn, can be more reliably used as an input to calculate the MVPF of the said inheritance tax reform, not unlike the procedure that unfolded in the past years for building the library of MVPF of top income tax rate reforms and MVPF of corporate taxes that are now in our disposal.

[19] See, e.g., (Löffler & Siegloch, 2021; Siegloch et al., 2021).

Chapter 3

Comparative Welfare Inferences Using MVPF Framework

Leveraging MVPF Results for Comparative Welfare Analyses and Policy Derivations

Chapter 3.1 Introduction

As hinted in Chapter 2, a complete welfare analysis that leverages the marginal value of public funds requires several further aspects to be shed light upon. First, when comparing MVPF values of two policies corresponding to two beneficiary groups, one actually needs to explicitly state the implicitly assumed ratio of social welfare weights between the two groups. Chapter 3 elaborates further on this crucial and arguably most policy-relevant, yet unfortunately up to now arguably less widely enunciated in the literature, aspect of using MVPF as keystones in a holistic welfare analysis framework.

The remaining of this chapter is organized as follows. Subchapter 3.2 delivers a brief yet holistic overview on the literature strand, onto which the following quantifications are best attached. This methodology is then applied in Subchapter 3.3 – first by applying the framework for the numerous reforms of public expenditure in Germany between 1990–2018, as reported by (Setio, 2021).Afterwards, a two-country comparative study is given by benchmarking with the collection of welfare-evaluated fiscal policies of the United States initially analyzed by Hendren and Sprung-Keyser (2020) and now commonly referenced in the public economics literature.

Chapter 3.2 Theoretical Background

In a broader sense, the marginal value of public funds continues the tradition in the field of public economics in evaluating the welfare changes brought about by a fiscal policy through some form of benefit-cost ratio. In an early work, even before the public finance literature took its modern, more quantitatively-oriented form with pioneering works by Samuelson and Musgrave in the 1950s, Pigou (1920, p. 11) pointed out as a central goal of welfare economics the study of "certain important groups of causes that affect economic welfare in actual modern societies". While proxies such as consumer surplus have generally been accepted for evaluating individual economic welfare, measuring aggregate welfare entails several extra layers of analysis (cf. Slesnick, 1998) – chief among them the specification of the planner's preferences for redistribution. The (interpersonal) welfare weights discussed in detail in Chapter 3 reflect one major approach to this aim.

To facilitate the cross-sectional comparisons using MVPF, I briefly highlight key aspects of the framework. The marginal value of public funds (MVPF) is defined as the ratio of aggregate willingness to pay (of the policy's beneficiaries) to the governmental net cost; it can be equivalently interpreted as the shadow price of raising revenue from the corresponding beneficiaries of the program expenditure; lastly but crucially, its design allows it to incorporate causal effects of policy changes obtained from increasingly rigorous toolbox of the empirical economics. In other words, the MVPF framework requires the net cost to the government to incorporate, on top of the actual government expenditure on the policy (mechanical cost), all behavioral responses from the beneficiaries (externalities) that are induced by the policy. These externalities are in turn estimated through the increasingly available and ideally causal empirical findings.

$$\mathrm{MVPF} = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to the Government}}$$

Estimation of willingness to pay follows Hendren and Sprung-Keyser (2020) in adopting the envelope theorem, where a marginal policy change exerts no immediate impact on marginal beneficiaries' individual utility functions. A marginal beneficiary is defined as an individual who creates behavioral response to the marginal policy change, i.e., one who was not part of the program *prior* but changes their behavior to become an eligible recipient *after* the marginal policy change (Finkelstein & Hendren, 2020). An infra-marginal beneficiary, on the other hand, is one who had been part of the program even before the marginal policy change. In the case of cash transfer, a marginal transfer of 1 unit is thus valued by its incumbent recipients, i.e., the infra-marginal beneficiaries, at exactly 1 unit (it is for them a straightforward cash transfer, after all). In other words, the only welfare effects such marginal policy change produces—the willingness to pay—is sufficiently captured by 1 unit times the number of infra-marginal recipients.

More relevantly to our understanding on MVPF's relation to and advantage against other welfare methods, I now turn once again to the historical review of welfare analysis. In his seminal work introduced earlier, Pigou (1920) acknowledged the inevitability of considering social welfare weights, stating, "If income is transferred from rich persons to poor persons, the proportion in which different sorts of goods and services are provided will be changed" (p. 89). This crucial notion of incidence-based externalities remains just as relevant in today's public economics and sets here the stage for a specific MVPF advantage against other welfare methods. To be specific, the MVPF as welfare method is most related in the public finance literature to marginal cost of public funds and marginal excess burden approaches.

In the *marginal cost of public funds* framework, one calculates the benefits of a policy and divides it by the costs needed to incur government revenue the size of the initial expenditure. These costs, in turn, consists of the actual spending for the policy and a proportional cost premium. The latter is known in the literature as the distortionary costs of taxation and is commonly assumed to be between 30–50% (see, e.g., Saez et al., 2012). By definition, depending on the group of individuals selected as program recipients, the true marginal cost of public funds will vary. The MVPF approach, on the other hand, can facilitate direct comparisons of policies across different recipient groups, due to its inclusion of social welfare weights in its calculation. It begins by defining aggregate social welfare as the weighted sum of individual utilities, whereby each of the weights corresponds to a specific individual and reflects the impact of a marginal increase to that individual's utility on aggregate social welfare. Multiplying this weight with the corresponding individual's marginal utility of income gives us η_i, which denotes individual i's social marginal utility of income, i.e., the impact of a marginal increase of individual i's income on aggregate social welfare. We further let $\bar{\eta}_j$ denote the average social marginal utility to policy j's beneficiaries, which consists of individual social marginal utility of income η_i and individual willingness to pay for the policy. In other words, the average social marginal utility $\bar{\eta}_j$ reflects the impact of a marginal change in policy j on aggregate social welfare. $MVPF_j$ is defined as the ratio of aggregate willingness to pay of the policy j's beneficiaries to government's net cost. Finally, the impact of a marginal increase in public spending for policy j on aggregate social welfare is given by $\bar{\eta}_j \times MVPF_j$ and this is where MVPF's defining characteristic teased earlier comes into play: Unlike older welfare methods such as the marginal excess burden method, no hypothetical lump-sum compensation is needed in the MVPF framework to close the budget constraint because a direct comparison of two policies forms a hypothetical budget-neutral policy change, i.e.,

$$\bar{\eta}_1 \times \mathrm{MVPF}_1 \lesseqgtr \bar{\eta}_2 \times \mathrm{MVPF}_2$$

Equivalently, an MVPF trade-off between two policies considers the ratio of average social marginal utility to each policy's respective group of beneficiaries, i.e.,

$$\frac{\bar{\eta}_1}{\bar{\eta}_2} \lesseqgtr \frac{\mathrm{MVPF}_2}{\mathrm{MVPF}_1}$$

If policy 1 has an MVPF of 1 and policy 2 has an MVPF of 2, then one prefers a marginal increase in policy 1 funded through a marginal decrease in policy 2 if and only if $\frac{\bar{\eta}_1}{\bar{\eta}_2} > \frac{MVPF_2}{MVPF_1}$, i.e., when one values providing 1 unit to beneficiaries of policy 1 more than providing 2 units to beneficiaries of policy 2. In a nutshell: In the MVPF framework, incidence always matters when comparing policies with different beneficiary groups.

Another strand of optimal taxation literature concerns mainly the *marginal deadweight loss*—or also known as marginal excess burden—of a policy. While the marginal cost of public funds measures the welfare cost of exacting tax on beneficiaries, the marginal deadweight loss accounts for the expected amount of additional government revenue through replacing distortionary taxes with a certain form of lump-sum compensation (Auerbach & Hines, 2002). In reality, however, lump-sum transfers are a rather rare policy instrument. Moreover, calculating the optimal size of these Hicksian compensating variations is an extremely challenging empirical task because they are entangled to individual consumer utility functions—which are primarily private information. The MVPF approach alleviates this problem by not having to close the budget constraint through hypothetical lump-sum taxes, forming hypothetical budget-neutral policies through directly comparing MVPF of two different policies (Finkelstein & Hendren, 2020).

Chapter 3.3 Main Results and Discussions

Comparing MVPF within-group, across reforms: Top Income Tax Rates

The first use of MVPF values is to juxtapose the results concerning different reforms very comparable beneficiary group. To begin, this example illustrates this function by focusing on one segment of the income distribution, namely the top 1% of earners along Germany's income distribution, which have been subject to several reforms on their corresponding tax income rates. Table 3.1 below, an adapted excerpt of Table 3 in Setio (2021), presents the calculated MVPF results that are then juxtaposed in the following discussion and graphical representation.

Policy	Cited Study(s)	MVPF
Top taxes		
Top tax 1990	Schellhorn & Gottfried (2004)	∞
Top tax 2004	Gottfried & Witczak (2009)	2.77
Top tax 2004	Schmidt & Müller (2012)	3.25
Top tax 2004	Werdt (2015)	2.83
Top tax 2004, average		3.83
Top tax 2005	Schmidt & Müller (2012)	2.54
Top tax 2005	Werdt (2015)	2.31
Top tax 2005	Doerrenberg et al. (2017)	3.40
Top tax 2005, average		3.14
Top tax 2007	Doerrenberg et al. (2017)	3.40

Table 3.1 MVPF results for selected reforms of Germany's top income tax rate

The 1990 top marginal tax rate reform yields a fiscal externality estimate with a negative, larger than 1 value, which results in a negative denominator of the MVPF and subsequently infinite MVPF, by definition. As noted by Hendren and Sprung-Keyser (2020), an infinite MVPF value in the context of marginal tax rate reform indicates that the tax rate prior to the reform lied "on the wrong side of the

Laffer curve". This argument would further support the suitability of the tax reform.

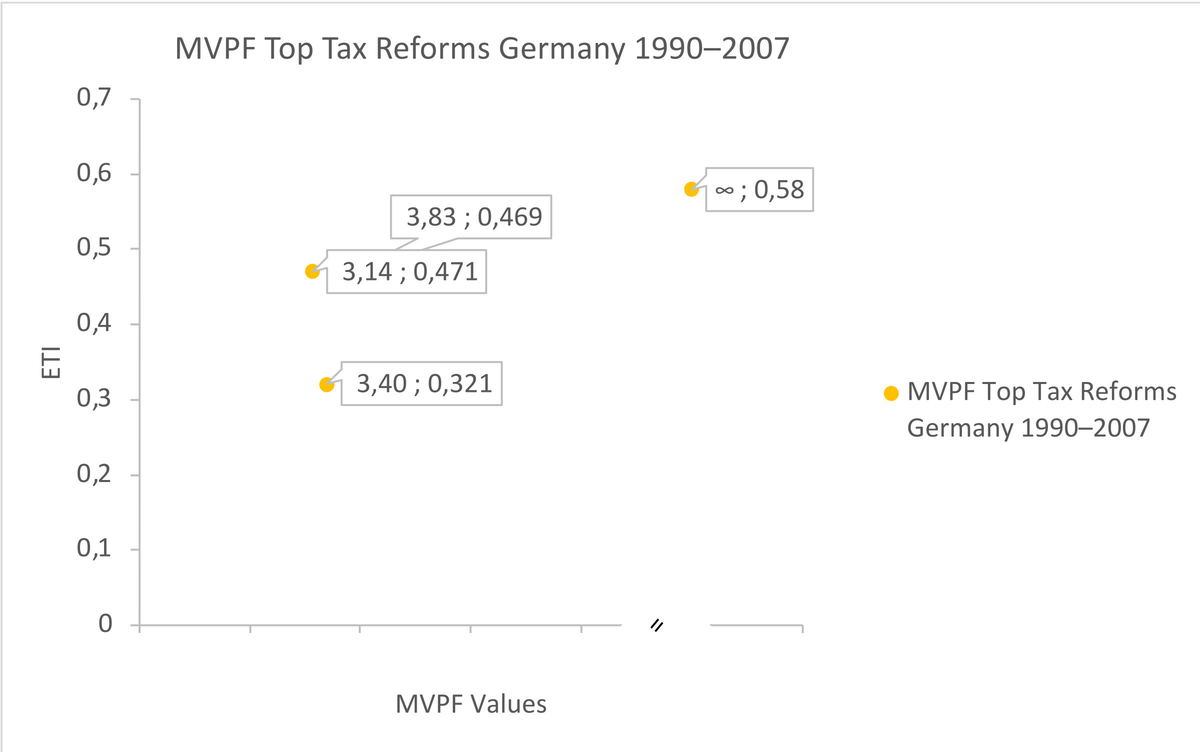

Figure 3.1 Visualizing MVPF results for selected reforms of Germany's top income tax rate. The vertical axis delineates the variations of the estimated elasticity of taxable income across years, despite focusing the analysis on one country. The labels are organized as {MVPF Value ; ETI estimate}.

None of the three subsequent tax reforms was found to generate negative net cost to the government, but all of them nevertheless yield positive, larger than 1 MVPF values. Extending positive welfare verdict to the 2004, 2005 and 2007 reforms would therefore be justifiable. However, it is also worth noting that for each estimation strategy the reforms appear to display a trend of decreasing MVPF over time. For example, based on Schmidt and Müller's (2012) estimate for elasticity of taxable income among married German households, MVPF of the 2005 reform (3.25) was lower than its 2004 counterpart (2.54).

A notable exception to this trend is the 2007 reform, for which the only available ETI estimate was that of Doerrenberg et al.'s (2017), whose shorter-term estimation strategy yielded consistently higher estimates[20]. A plausible explanation for this anomaly can be derived from Schmidt and Müller's (2012) robustness check, in which they found greater elasticity from using 1- and 2-year-difference specifications instead of their preferred specification's 3-year difference. Schmidt and Müller further noted that this was in line with the larger ETI estimates (0.58 and 0.44, respectively) of their two immediate predecessors (Gottfried & Witczak, 2009; Schellhorn & Gottfried, 2004) who both employed 2-year difference. ETI estimates with 3-year difference model align better with Saez et al.'s (2012) summarizing literature, which concluded that plausible estimates should be within the range of 0.12–0.40. In other words, current evidence would seem to caution against using too compact time period between the pre- and post-reform observations. Schmidt and Müller argue for short-term behavioral responses such as income shifting within adjacent years as one of the plausible reasons why shorter time lags appear to yield greater elasticity results.

Interestingly, Hendren and Sprung-Keyser (2020) documented a similarly downward trend for top marginal tax reforms in the United States during 1981–2013. There, the top marginal tax reform in 1981 resulted in an infinite MVPF, whereas the most recent reforms in 2001 and 2013 yielded MVPF of 1.37 and 1.16, respectively.

One caveat of the MVPF approach in the taxation context that becomes clearer through this exercise of implementing MVPF framework for the German top marginal tax rate reforms is the inability of MVPF framework to distinguish the direction of behavioral response between a tax cut and a tax hike. Because the final fiscal externality is calculated as an average of the fiscal externality of the tax rate *before* reform and *after* reform, and because in most cases both of these fiscal externalities turn out to be of negative values, the averaged fiscal externality is subsequently almost always negative. An obvious observation would be the comparison between the 2005 tax rate cut of 45% to 42% and the 2007 additional

[20] MVPFs for 2004 and 2005 reforms based on their ETI estimate are also considerably higher than the ones based on Schmidt and Müller's (2012) or Werdt's (2015).

introduction of top marginal tax rate at 45% using ETI estimates by Doerrenberg et al. (2017), which resulted in the exact same MVPF values. It is an important avenue of future research to develop the MVPF framework further in the direction of making it capable to distinguish the effects of opposite directions of tax rate reforms.

Comparing MVPF within-sector, across reforms: Job Training Policies

Figure 3.2 visualizes MVPF results for representative reforms of Germany's active labor market policies. This technique of comparison highlights the chief benefit of MVPF: facilitating direct, same-metric comparisons across previously separate policies. The results emphasize the potential efficiency differences across choices of programs, despite focusing the analysis on one beneficiary group (namely those at the participation margin of the labor market). Details on the depicted job training policies are re-iterated below.

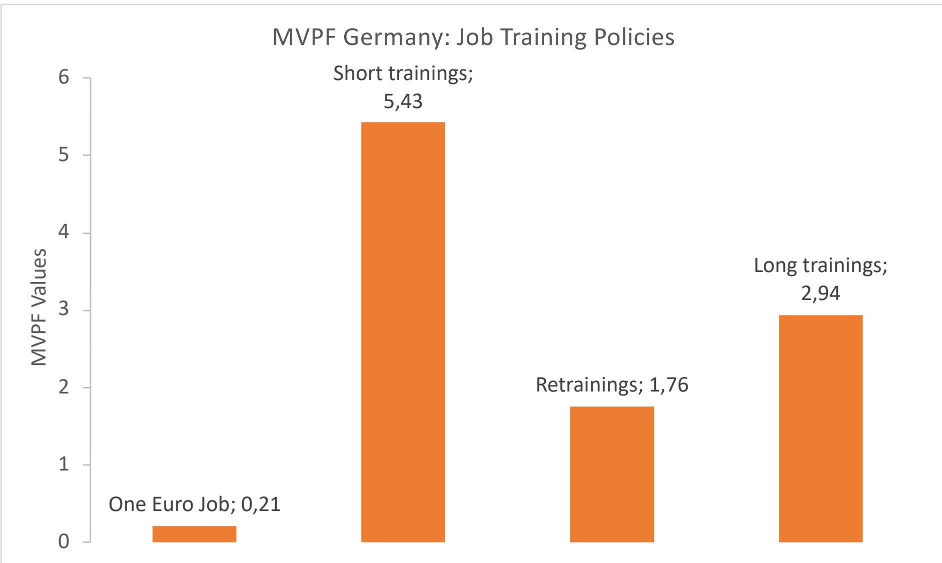

<u>Figure 3.2</u> Visualizing MVPF results for representative reforms of Germany's active labor market policies. The results emphasize the potential efficiency differences across choices of programs, despite focusing the analysis on one beneficiary group.

Ein-Euro-Job

The One Euro Job program (*Ein-Euro-Job)* was a job assistance first implemented by the German government in 2005–2007. The program offers auxiliary job for recipients of unemployment benefit (*Arbeitslosengeld II)* with a most dismal prospect of re-entering workforce. Harrer and Stockinger (2019) evaluated the effects of One Euro Job program after its reform in 2012. The One Euro Job program had an average monthly cost per participant of €459. The average participation duration was 4 and a half months. Willingness to pay for One Euro

Job program is derived from the received work compensation (€124 monthly on average) and change in earnings. Below the corresponding MVPF calculation.

Evaluating One Euro Job program at the end of the 4.5-month average participation duration and assuming worst case scenario of earnings loss as reported by Harrer and Stockinger (2019) yields the MVPF value of 0,21[21]

Short-term trainings

Bernhard's (2016) investigation documented similarly short lock-in effect for short-term training programs with participants experiencing on average €274 reduction to their monthly income for the first 6 months but benefiting from €286 more income per month for the rest of the 104-month-long observation period. I use Lechner et al.'s (2011) cost estimate (see Subchapter 2.3.1 for more details) and align the observation period with theirs (8 years) for more accurate comparability, arriving at

$$\text{MVPF} = \frac{-€274 \times 6 \ + \ €286 \times 90}{€4{,}439} = 5.428$$

Long-term and (re-) trainings

Calculations for the two remaining types of job training the selected sample of German ALMP policies proceed analogously. To illustrate briefly, using *long* training cost per participant €9,930 instead affects the calculation as follows.

$$\text{MVPF} = \frac{-€340 \times 24 \ + \ €416 \times 72}{€9{,}930} = 2.944$$

[21] The general formula for programs exclusively targeting labor force at participation margin can be found in Subchapter 2.3.1.

Figure 3.3 *Visualizing MVPF results for representative reforms of Germany's active labor market policies. The results emphasize the potential efficiency differences across choices of programs, despite focusing the analysis on one beneficiary group.*

To reinstate the main argument of Chapter 3, drawing welfare conclusions based on head-to-head ranking of MVPFs across domains, it is necessary to keep in mind that this framework inherently implies quantification of intergroup trade-offs. To see this mechanism in concrete example, recall the stark MVPF contrast between One Euro Job and other types of public job training. One Euro Job lasts on average 4.5 months (Harrer & Stockinger, 2019), which is a training period most comparable Lechner et al.'s (2011) definition of short training. The former has an MVPF of 0.21 and the latter 5.43, which is over 25 times higher. Strictly speaking, the policymaker prefers a marginal increase in expenditure for One Euro Job funded through a marginal decrease in expenditure for short training if and only if she values providing 0.21 monetary unit to One Euro Job beneficiaries more than providing 5.43 monetary unit to beneficiaries of the competitor program, i.e., if and only if $\frac{\bar{\eta}_{One\ Euro\ Job}}{\bar{\eta}_{Short\ training}} > \frac{\text{MVPF}_{\text{Short training}}}{\text{MVPF}_{\text{One Euro Job}}} = \frac{5.43}{0.21} = 25.86$. Here, the left-hand side of the inequality is the crucial takeaway. Without keeping in mind the necessary consideration of social incidence when comparing MVPFs, one risks running into false, or at least premature, conclusion. Here, a background check on the designated beneficiaries of each program would serve us well: the One Euro Job was designed for unemployed individuals with (under certain criteria) the most desperate outlook for self-reintegration into the labor market, i.e., as an *ultima ratio*. On the other hand, the short trainings per Lechner et al.'s (2011) definition are not strictly confined to individuals most-in-need.

Comparing MVPF across sectors, within-region: Social Insurances

Figure 3.4 continues the chapter by visualizing MVPF results for representative reforms of Germany's social insurance policies. This comparison highlights another main advantage of MVPF, namely, enabling unified juxtapositions across sectors of social insurance policies, each of which previously had its own welfare metric (such as the Baily-Chetty condition for health insurance; the BC/MC ratio for unemployment and disability insurances, on the other hand).

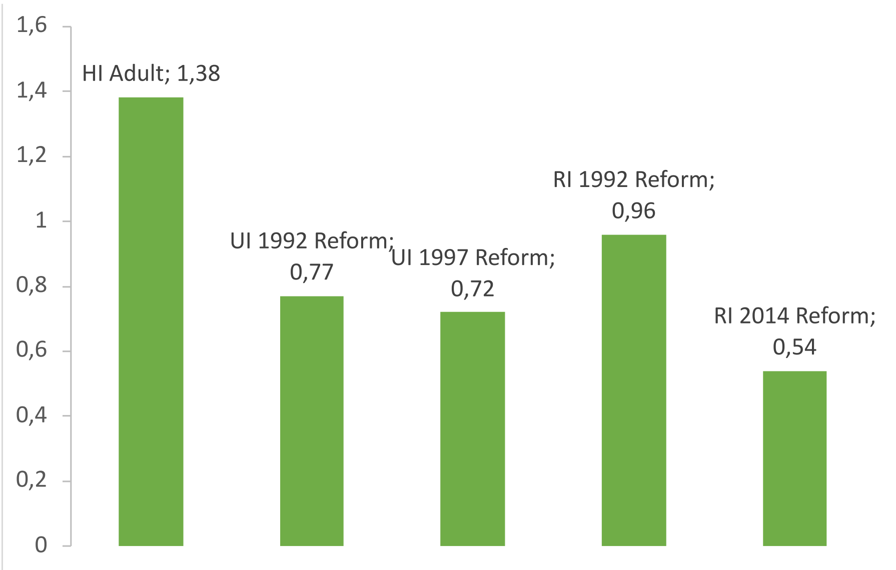

Figure 3.3 Visualizing MVPF results for representative reforms of Germany's social insurance sectors.

Health insurance

Using 2011–2014 LIFE-Adult-Study dataset from Leipzig Research Centre for Civilization Diseases, Hajek et al. (2020) reports a monthly average willingness to pay of €240 for health insurance, which represents approximately 14% of average net monthly household income in Germany.

Federal Statistical Office's official health system databank Federal Health Monitoring (Gesundheitsberichterstattung des Bundes, 2021) documented an annual average government expenditure per public-health-insured individual of

€2,057 in 2010, €2,090 in 2011 and €2,355 in 2014. I combine each of the latter two with Hajek et al.'s (2020) willingness-to-pay estimate to form an MVPF_{2011} of $\frac{€240 \times 12}{€2{,}090} = 1.378$, $\text{MVPF}_{2014} = \frac{€240 \times 12}{€2{,}355} = 1.223$, and $\text{MVPF}_{\text{Average 2011–2014}} = \frac{€240 \times 12}{€2,} = 1.296$.

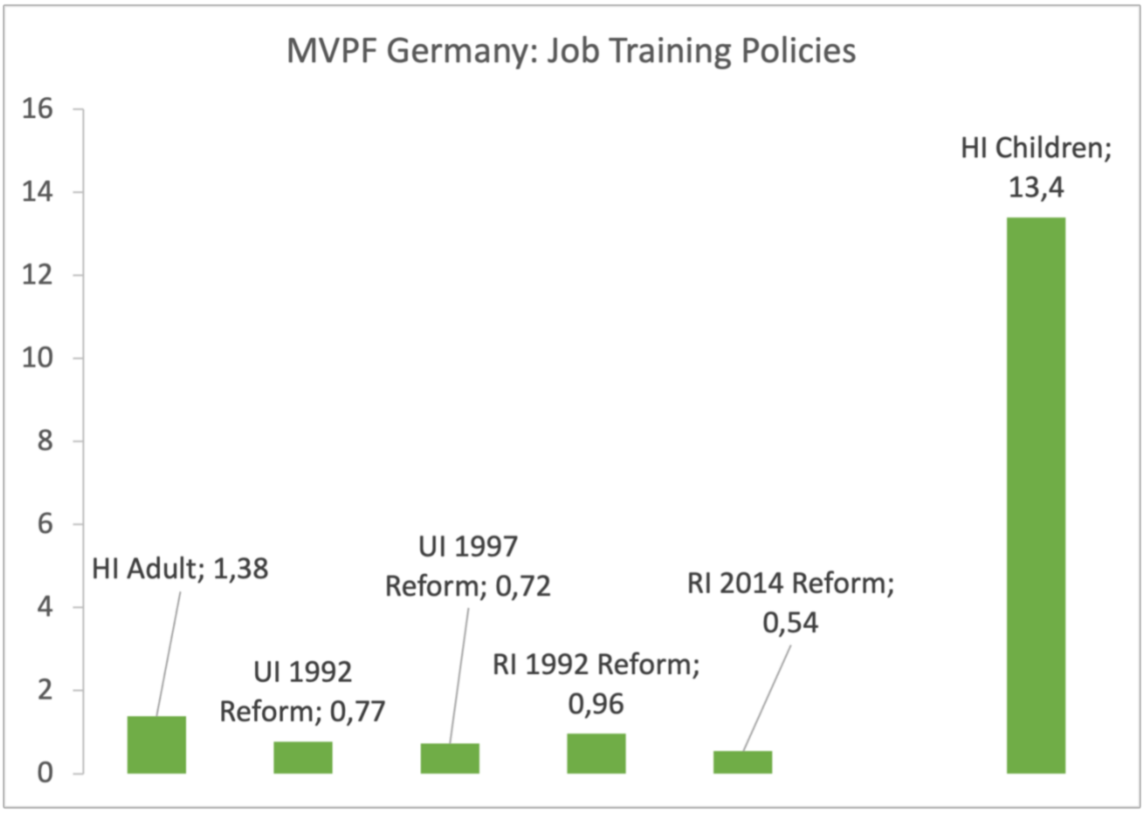

Figure 3.4 *Graphical comparison of various German social insurance polices. Including reform Join the Healthy Boat, which targets specificially children under 14 years old as beneficiary group.*

Unemployment insurance

Several authors (e.g., Schmieder & von Wachter, 2017; Ye, 2018) examined the ratio between behavioral and mechanical costs (henceforth BC, MC) induced by an unemployment insurance policy. The rationale for adopting BC/MC ratio in unemployment insurance context can be compared to the motives behind the success of elasticity of taxable income (ETI) as a measure of fiscal externalities due to a given tax change. ETI captures behavioral responses to a tax reduction/increase in a more comprehensive way than the traditional labor supply elasticity.

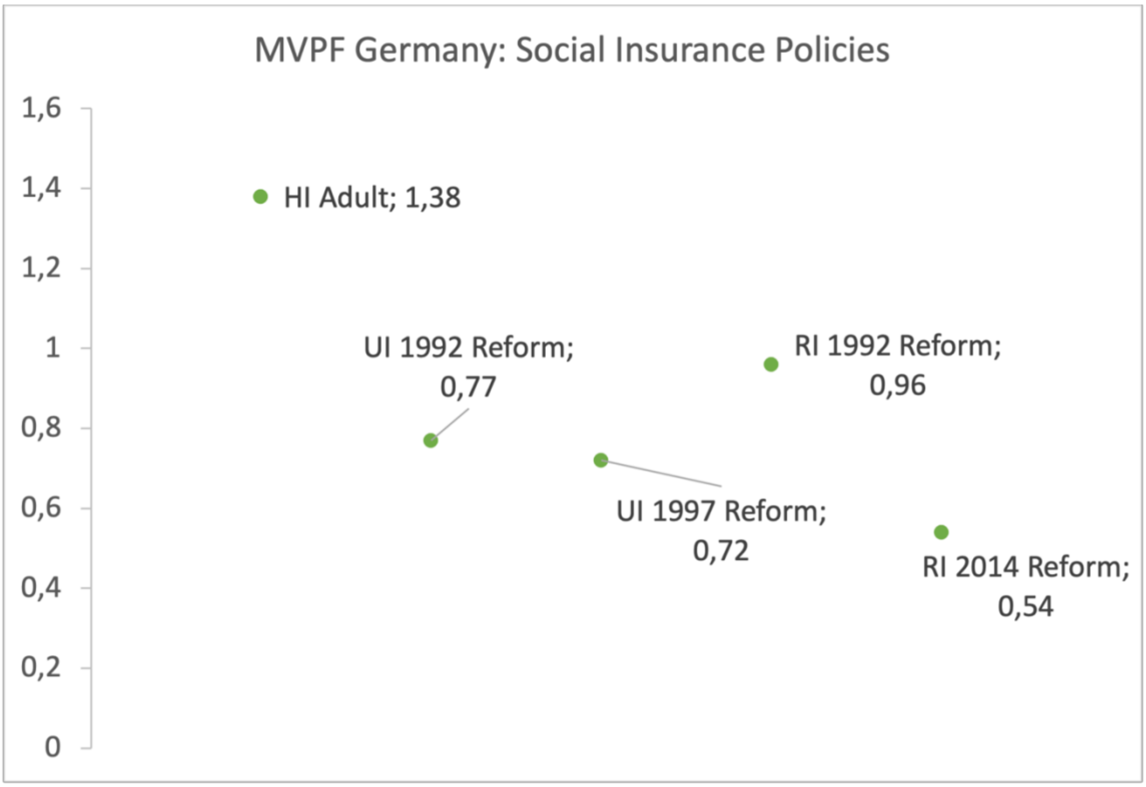

<u>Figure 3.5</u> Visualizing MVPF results for representative reforms of Germany's social insurance sectors.

Translating into the MVPF framework, mechanical cost can be interpreted as the amount of expenditure the government spends for funding the program, whereas behavioral cost captures additional costs to the government due to behavioral responses of the beneficiaries toward the policy, a concept similar to fiscal externality. Recall the earlier established notion that for a marginal policy change, willingness to pay is defined as equivalent to mechanical cost, assuming the envelope theorem holds (following this assumption, any change to the individual's budget constraint is attributable to the policy change, and not to changes in labor effort). In other words, the MVPF can be equivalently defined as

$$\text{MVPF} = \frac{\text{Mechanical cost}}{\text{Mechanical cost} + \text{Behavioral cost}}$$

Normalizing for mechanical cost gives an equivalent definition as follows

$$\text{MVPF} = \frac{1}{1 + \frac{\text{Behavioral cost}}{\text{Mechanical cost}}}$$

In 1992, Germany increased pension benefits for low-income workers as part of the Pension Reform Act. Although at first glance this policy may seem to completely fit into the landscape retirement insurance, a closer look at retirement entry schemes of older German workers reveals another well-trodden path of entering retirement through periods of unemployment. The behavioral response to the policy, which is crucial to the MVPF concept, must therefore be incorporated in the efforts of mitigating excessive unemployment level. Ye (2018) evaluated this 1992 reform and found a behavioral-to-mechanical cost ratio of 0.3. Translating this estimate into MVPF thus gives $\text{MVPF} = \frac{1}{1 + 0.3} = 0.769$.

Comparing MVPF within-sector, across-regions

Job Training Policies in Germany and United States[22]

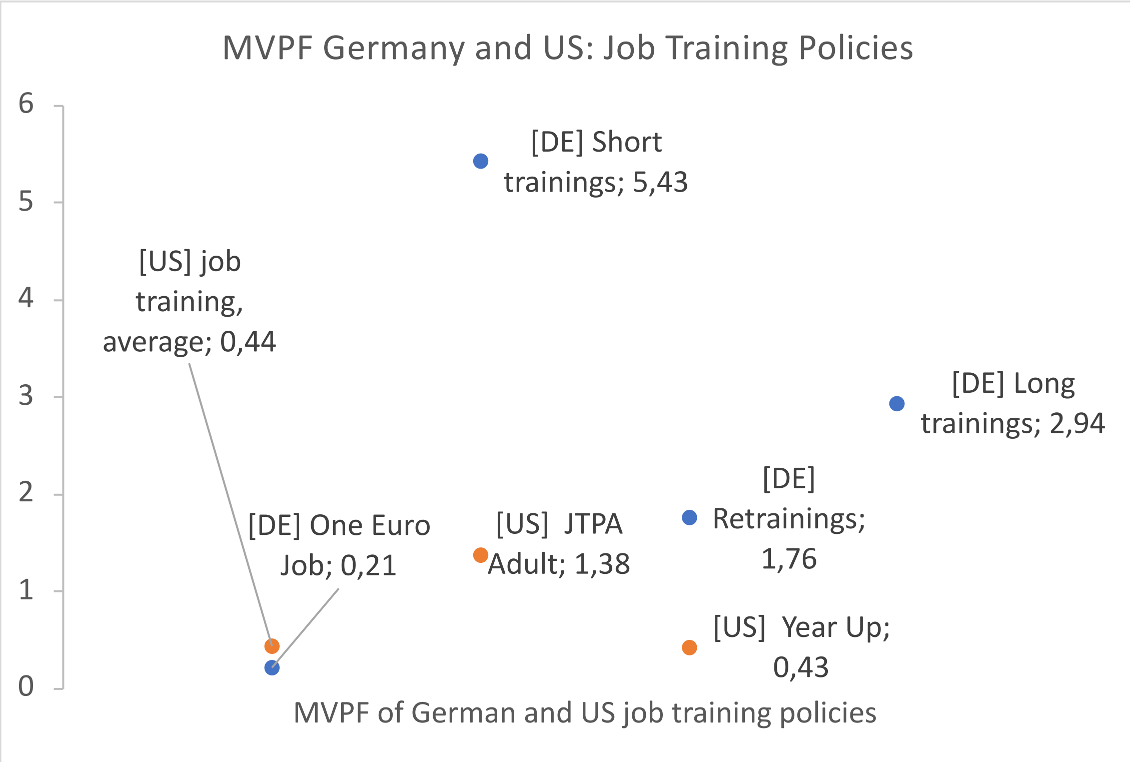

Figure 3.6 *Comparing selected MVPF results of Germany's vis-à-vis United States' programs in job training policies.*

For publicly administered job training programs, short-term trainings (less than 6 months) tend to display significantly higher MVPF than longer-term trainings. It is worth to take note that my calculations in this regard inevitably depended on the available timeframe of earnings projection in the source studies. For example, it could be argued that since retraining programs as defined by Lechner et al. (2011) award their participants upon completion a new professional degree, their benefits in form of income surplus will continue to be reaped years into the future, well beyond the observation period currently available—eventually inflating the program's MVPF. The relative ineffectiveness of the One-Euro-Job program, on the other hand, is quite unambiguous since their estimated MVPFs are a lot

[22] Calculations on these US MVPFs are given in Hendren and Sprung-Keyser (2020).

smaller than other programs, even smaller than 1. This would appear to concur with the findings of Vooren et al. (2019) that public employment/job creation generally create more persistent negative employment effects than other program types such as subsidized labor or job search assistance.

Comparing MVPF across-sector, across-regions

Social Insurance Policies of Germany's vis-à-vis US'

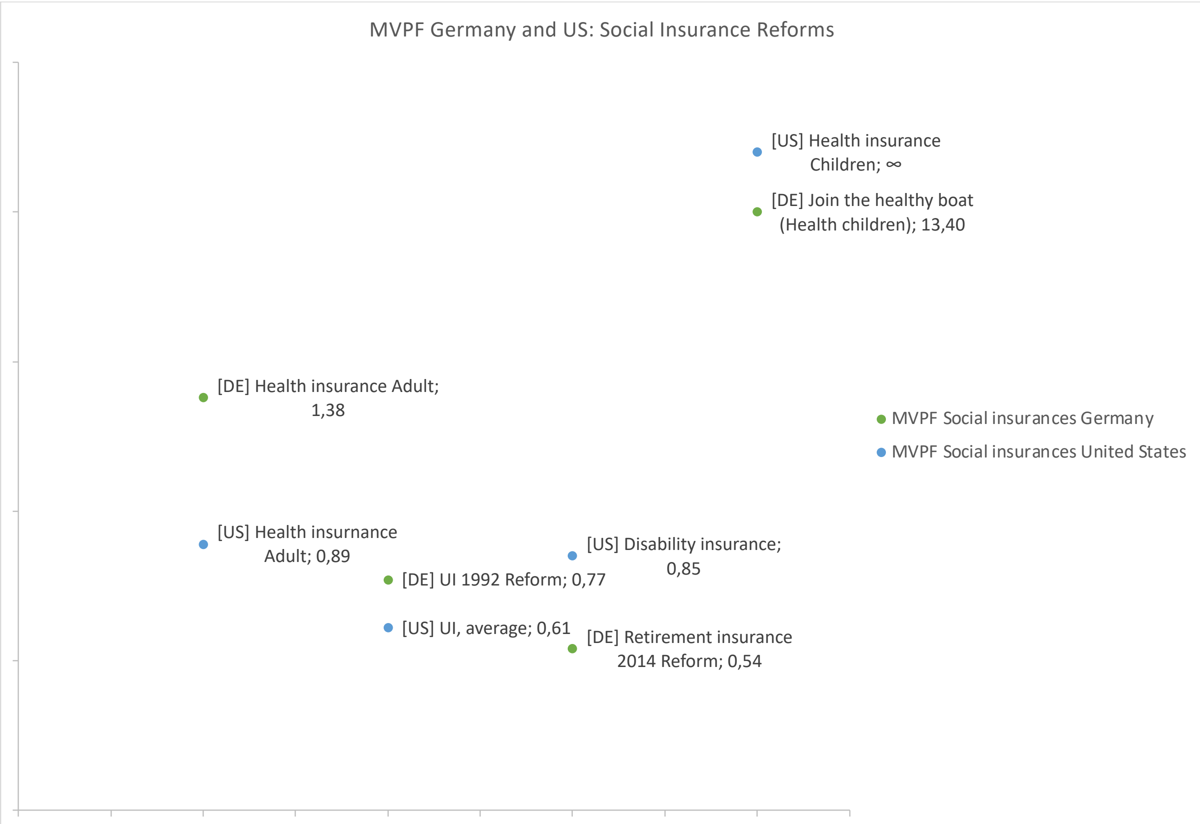

Figure 3.7 Comparing selected MVPF results of Germany's vis-à-vis United States' programs across social insurance sectors.

MVPF results for German unemployment insurance display a comparatively limited amount of variation between themselves, averaging 0.71 with standard deviation of 0.043. These results appear aligned with findings for unemployment insurance policies in the United States which had on average MVPF of 0.61 without large variations.

MVPF results for German public health insurance imply positive welfare effect, albeit at around 1.30 a relatively small one. Slightly higher MVPF (1.52) for senior citizens reflect this cohort's higher willingness to pay as discovered by Bock et al.

(2016). These figures are remarkably similar to Hendren and Sprung-Keyser (2020)'s MVPF findings for United States' health insurance policies, the highest of which was 1.63. Also much in line with each other are MVPF findings for the two countries' health insurance policies whose primary concern are children—out of four US child health policies, the lowest MVPF found was 10.24, which is fairly close to the average MVPF of German school-based campaign Join the Healthy Boat (12.22).

Chapter 3.4 Concluding Remarks

A common caveat across empirical studies on the welfare impact of social insurance programs is the varying extent of externalities coverage and their inconsistent levels of budget disaggregation can be readily apprehended by observing policies of social insurance. Within one domain, treatment intensity for fiscal externalities differs contrastingly across existing empirical studies. While the behavioral/mechanical cost approach common to unemployment insurance evaluations by definition takes into account all behavioral responses and recent authors such as Krolage (2020) incorporated fiscal externalities in her estimate of total costs of pension reform, no source has documented how German preventive health programs might lead to less expenditure on remedial measures or to higher tax revenue due to increased productivity of healthier workers.

In conclusion, Chapter 3 highlighted possibilities to utilize a given collection of MVPF results. Discussions on these aspects each draws on dynamic discourse in the literature, representing on the one hand several of the common commentaries, while also being pointed out as key differences to previous literature. In general, welfare comparisons that utilize MVPF can be carried out along the dimension of: 1) within-group, across reforms; 2) within-sector, across-reforms; 3) within-region, across-sectors; and 4) country-comparative study, within-sector.

Chapter 4

Policy Redistribution Using Statistical Decision Theory

Theoretical Summary and Methodological Extensions of Statistical Treatment Allocation for Fiscal Policies

Chapter 4.1 Introduction

In Chapters 2 and 3, the problem of heterogeneity across public policy domains has been introduced and the methodology of Marginal Value of Public Funds was utilized as a unifying welfare analysis framework. The second half of this dissertation turns to the problem of individual heterogeneity that causes the expected treatment impact of any given policy allocation to be heterogenous and hinders subsequently calculation of optimal policy from being straightforward. At this point, the policy-making process has progressed from taking the bird's eye view – learning and ranking marginal values of public fund across policy fields – to specifying the expected properties of a given policy field. This paradigm is reminiscent to the hierarchical framework of optimization levels found in related literatures on decision sciences. From this perspective, building policy priorities based on ranked MVPFs (using contributions such as in Chapters 2 and 3 of this dissertation) could be seen as representing a long-term, strategic perspective, whereas ensuring the chosen policy generates optimal welfare effect via the right choice of beneficiaries (the topic of Chapters 4 and 5) belongs to medium-term, tactical approach.

One suggested avenue suggested by (Hendren & Sprung-Keyser, 2020, p. 1222) for further research is the potential integration of effect heterogeneity into MVPF framework. In particular, within a given policy context for which there already exists some established MVPF evidence, the policymaker faced with different budget size and allocation would need to depend on some assumptions, unless she investigates further the average treatment effect for each policy configuration. Her task gains another layer of complexity when the population considered is heterogenous with regards to individual willingness-to-pay for the treatment.

In some cases, accounting for heterogeneity can open up new line of evidence where the literature had previously followed an opposite consensus. For example, Caliendo, Hujer and Thomsen (2008) showed that while prior studies found negative average treatment effect for German job creation schemes (JCS), disaggregating enabled them to present a clearer picture of the situation. In other words, they discovered no negative effects for most segments yet singularly positive effect for the long-term unemployed.

Impetus towards identifying optimal treatment assignment rule for a given policy already existed in several previous studies (e.g., Lechner et al., 2011 for active labor market program assignments). Such an attempt was usually presented in auxiliary section on sensitivity analysis, as a means to strengthen the study's main claim. The choice of alternative assignment rules, despite accompanied with adequate justification, were less than perfectly systematic. On the other hand, the Empirical Welfare Maximization method (Kitagawa & Tetenov, 2018) is appropriate to advise treatment assignment for any policy, hence its adoption would reduce the time and effort otherwise needed in building and evaluating alternative scenarios.

The incorporation of advanced statistical tools (and recently, even that of machine learning) toward optimal public policymaking gained significant momentum through the seminal contribution by Manski (2004), who first identified treatment allocation problem as a statistical decision problem. In other words, considered is a policymaker who observes covariates of all population members *and* knows the sample's response to the treatment, *but* lacks information on the population distribution of treatment response. In this case, projecting sample treatment response into population requires some statistical treatment rule. This line of research was carried forward by Dehejia (2005), Hirano & Porter (2009), Stoye (2012), and most recently, Kitagawa & Tetenov (2018) and Athey & Wager (2021).

The Empirical Welfare Maximization method minimizes the welfare loss caused by the heterogeneity of the covariates, which can be described as a negative deviation from the ideal (first-best) welfare potential, especially in the context of policy measures due to social-political conditions. Kitagawa and Tetenov (2018) performed an application of this modern method using an empirical example in the context of the United States and achieved up to $897 additional welfare premium per individual compared to the manual calculation of the National Job Training Partnership Act (JTPA) program implemented in 1982. In total, the use of the EWM method in JTPA program would have resulted in a benefit of over $6 million (75% treated, sample size 9,223).

Athey and Wager (2021) have developed a similar method that specifically introduces recent findings from the machine learning field into the context of optimal policy measures. An R package called policytree was published as implementation support. Unlike Kitagawa and Tetenov (2018), Athey and Wager focused on the application of statistical decision trees (commonly known as random forests among data science disciplines) as the main tool for determining optimal policy distribution. The main advantage of their policy learning method compared to Kitagawa and Tetenov's is its ability to learn from observational data, i.e., no costly/infeasible field experiments have to be performed beforehand in order to supply the machine learning algorithm with data. This represents an excellent potential due to the increasingly available administrative big data across government institutions.

Chapter 4.2 Literature Review

This chapter reviews, applies, and extends recently proposed methods for learning optimal treatment assignment policies. Starting with Manski (2004), this literature combines semi- and non-parametric methods to derive statistical treatment rules that maximize welfare, i.e., policy gain. The current state of the literature culminated with Kitagawa and Tetenov's (2018) *Empirical Welfare Maximization*, Athey and Wager (2021)'s *Policy Learning*, Mbakop and Tabord-Meehan (2021)'s *Penalized Welfare Maximization* and Manski (2021)'s *Asymptotic Minimax Regret*. This chapter benchmarks these four methods against each other methodologically. It further contributes to the literature by introducing several methodological extensions: a) modifying the methods with alternative advanced machine learning *penalized least squares* neural networks, *boosting and support vector machine* and compare their performance; and b) supplementing the methods with techniques from the growing subfield of interpretable machine learning. The chapter equally serves to illuminate the theoretical underpinnings of the complete statistical method that will be utilized for empirical applications in Chapter 5, which juxtaposes the empirical performances of both the status-quo (Policy Learning), the aforementioned benchmark methods corresponding to it (Empirical Welfare Maximization, AMMR and PWM), as well as my own methodological developments on two large-scale administrative datasets: Germany's SIAB and Indonesia's JKN-KIS[23].

Heterogeneity within observed data can in principle be exploited to improve policy evaluation and design, such as by inherently incorporating it into treatment allocation procedure. An early methodological development in this area was to account for local average treatment effects (LATEs) – for discussion see, e.g., Athey and Imbens, (2017). Other advances include refinement methods that aim to identify how variation in individual treatment effects affect causal estimate of a parameter, culminating in Athey, Tibshirani and Wager (2019).

[23] JKN-KIS is the Indonesian dataset for the publicly-administered health insurance program.

While considerable emphasis within the econometric literature has been given towards incorporating heterogeneity into identification and inference toolkit, prediction problems increasingly benefit from the growing field of statistical learning and the fruitful intersection between econometrics and machine learning. Predictive models and learning methods have been employed in econometric applications across various fields, an example of which includes predicting university student dropout rate (Kemper, Vorhoff and Wigger, 2020)

Yet another type of task universally related to policymaking beside inference and prediction is treatment choice. To put it concisely, treatment choice seeks to illuminate the question of *whom* to treat with *which* treatment specification while exploiting individually heterogeneous treatment effects, respecting exogenous constraints such as budget but also political, legal or ethical constraints, and maximizing a desired welfare criterion (such as the minimax-regret (MMR), which is especially utilized in the empirical implementations in Chapter 5). It differs fundamentally from both prediction and inference problems in that it integrates the welfare maximization or regret[24] minimization directly into its optimization procedure, as opposed to the "plug-in" or "two-step" approach where the results of identification from previous stage are consequently used in a separate state. Treatment choice seeks to optimize across all possible states of nature given uncertainty with inherent weighting of the performance of the chosen treatment rule in each possible state, instead of purely optimizing the statistical distribution without regard towards the real economic incidence. Manski (2021) aptly summarizes this distinction by drawing a parallel to the old discourse between normative and positive economic theories: where prediction and inference as "descriptive" decision analysis are intended to explain and predict how decisions are actually made, statistical treatment choice as "prescriptive" analysis aims to deliver context-relevant and implementable policy recommendations, even when some deadweight loss due to some external constraints were unavoidable.

[24] In the treatment choice literature for public policies, „regret" is defined as the „risk" or the expected welfare loss of using the proposed treatment rule instead of a hypothetically optimal (but unattainable) treatment rule.

A relatively small but growing strand in the econometric literature initiated by Manski (2004) strives to adopt statistical decision theory into and demonstrate the advantages of its use for solving econometric problems, particularly in the context of public policymaking. A partial list of those studies include (Armstrong & Shen, 2014; Athey & Wager, 2021; Bhattacharya & Dupas, 2012; Dehejia, 2005; Hirano & Porter, 2009, 2020; Kasy, 2018; Kitagawa & Tetenov, 2018; Mbakop & Tabord-Meehan, 2021; and Stoye, 2009, 2012). The foundations of statistical decision theory itself was laid in Abraham Wald's seminal work *Statistical Decision Functions* (1950). The theory is primarily concerned with providing a decision-maker who is facing uncertainty in her assessment of possible actions with a prescription that leverages statistical information from an existing sample data. The desired outcome is thereby a systematically selected function mapping the data generated from some sampling distribution into an action. Such a function is called in this framework a statistical decision function[25].

In a recent recapitulating work, Manski (2021) argues for a universal adoption of statistical decision theory as a selection mechanism over relevant models in classic econometric problems such as treatment choice and prediction. In prediction models, the unknown states of nature consist of the distributions of a real random variable whereas the decision under consideration is prediction of a realization drawn from the true distribution. For treatment allocation, possible distributions of individual treatment response constitute the stochastic states of nature. Statistical decision functions in this setting are by convention called statistical treatment rules.

Optimal treatment allocation has also been under considerable research spotlight in several neighboring fields. In medicine and epidemiology, it is often called individualized treatment rules. In data science and machine learning, the

[25] A statistical decision function is ex post deterministic, i.e., after a corresponding sample has been generated. Before the sampling procedure, however, it is random. An important implication of this is that the outcome variable (e.g., for this context, achievable welfare) is consequently also an ex-ante random variable.

optimization procedure is known as the problem of learning treatment assignment out of individual characteristics, or shortly, policy learning. For thorough discussion on the links between these literatures see, e.g., Athey & Wager (2021)[26], and Mbakop & Tabord-Meehan (2021).

[26] Athey & Wager (2021) incorporate elements from the machine learning literature into their statistical optimization method for treatment allocation and accordingly name it *policy learning*.

Chapter 4.3 Benchmark Analysis of Current Methods

This section investigates common theoretical elements between Policy Learning and its nearest counterparts in the literature strand with particular focus on identifying key distinctions across the benchmark methods. In the following, the four methods are each described and explained, while being altogether organized in an order that particularly takes into account their respective methodological complexities.

Empirical Welfare Maximization (EWM)

Synthesizing theoretical advancements in the growing literature strand on statistical decision theory for policy assignment ground-laid by Manski (2004) described above, Kitagawa and Tetenov (2018) assembled the Empirical Welfare Maximization, short EWM, method for determining the portion/s of observed sample population that “should” (with the optimization function of obtaining optimal welfare in aggregate) be “treated” (assigned) the policy in-question. As later elaborated by subsequent studies who explore group average treatment effects (GATE) such as Knaus (2022), here Kitagawa and Tetenov (2018) already noted that the one of EWM method’s main advantages is that it bypasses the need to perform individual regression estimates for every individual treatment effect (a task computationally nigh-impractical in light of ever-growing size and complexity of available administrative datasets for public sector, some of which overgrow a hundred in their amount of covariates as potential independent variables, let alone the size of instances/individual observations that can run well over a few millions—two empirical settings in Chapter 5 give example to this development).

Formally, the Empirical Welfare Maximization method defines the objective function of optimal treatment assignment “rule” (assignment scheme/“whom-to-treat”, which is the goal of these methods and serves as policy recommendation) as being the product of expected individual outcome and that individual’s *inverse-probability-weighted* (IPW) propensity score. Depending on the empirical setting, expected individual outcome can translate into, e.g., expected quarterly wage premium post-job-market-training-assignment in the context of active labor market programs such as the one studied later in Chapter 5.3 of this dissertation; or expected additional efficiency gains from re-prioritization within a multi-stage referral system in a health insurance context (Chapter 5.4). The latter variable, the IPW propensity score, is in essence a *weighting* method assigned to each individual in order to adjust their impact on aggregate welfare estimation. Unlike older regression adjustment methods such as matching estimators, however, the IPW has the advantage of estimating the propensity score values *in a separate step* before subsequently using them in inverse proportions to re-balance the weights of individuals in the sample. Without this feature, the fact that EWM works with only one-off (non-repeated) type of datasets (a capability which is precisely the reason public economics datasets such as SIAB and BPJS-KIS *can* be analyzed with it) would have made it susceptible to *bias*, owing to the skewness caused by the over-proportional weighting of observations with high propensity score, who are by definition overrepresented in the treatment (ex-ante policy-assigned) group. In short, the use of IPW weighting method and its combination with individual sample treatment effect is the core methodological breakthrough achieved by EWM which renders it, unlike prior statistical learning methods solely suitable for experiment-based data types, also suitable for observational datasets like the ones public or labor economics literatures are provided with[27]

Penalized Welfare Maximization (PWM)

The Penalized Welfare Maximization (PWM) method builds directly on theoretical framework and especially the mathematical guarantees given proofs in EWM. Mbakop and Tabord-Meehan (2021) expands the feasible data requirements of EWM into flexible re-arrangements of classes (the dimensions of partitioning rules allowed—e.g., one-dimensional classes of rules includes dividing job training assignment solely based on *one category* such as gender or age; on the opposite, allowing for more-dimensional linear rules and other forms of rules such as

[27] For reasons of exposition, formal definitions and mathematical properties are referred directly to Kitagawa and Tetenov, 2018.

decision trees would allow for a *combination* of these covariates either additively or in subsequent fashion). The methodological breakthroughs of this method appear to pertain above all to empirical settings where the overall pool of available classes of policies supersede those found in settings of labor market programs or social insurance; a reason I chose to focus for my empirical applications in Chapter 5 instead on EWM and the third method discussed below: the Policy Learning method. I note here, nevertheless, that the Penalized Welfare Maximization method represents an equally important development in the literature since EWM, and one that is potentially even more relevant to future studies with empirical settings where the legal or ethical constraints necessitate the flexibility of the *types of policy assignment rules* to be expanded into non-binary forms.

Policy Learning

The Policy Learning method developed by Athey and Wager (2021) represents the culmination of this literature strand. Not only does the method improve upon the EWM and PWM with regards to the scope of empirical application feasible, but it also conjures up a confluence with the methodological progress achieved in the strand of statistics literature called *doubly robust analysis*, which achieved its own culmination in the seminal contribution by (Chernozhukov et al., 2018)[28]. In a nutshell, the core improvement pertaining to the aggregate welfare (recall that the objective function of all these methods' optimization is the utilitarian, additively aggregated welfare gains due to the policy assignment) achieved by Policy Learning compared to its predecessors is in and due to its use of *doubly robust estimator*. Doubly robust estimator, also abbreviatingly called *double machine learning* in the computer scientific machine learning literature, involves the inverse probability weighting (IPW) feature also contained by EWM, but adds a *second layer* of machine-learning-supported estimation of the whole data sample. In other words, by doing this Policy Learning *leverages* the statistical power of the given data sample *twofold* – once, in a sense equivalent to matching and

[28] This paper, in turn, represents a continuation several papers by the authors, who hail from various fields: economics, econometrics and statistics. The eclectic nature is subsequently appropriately reflected in the remarkably wide acceptance and practice of *doubly robust estimator* across fields as diverse as biotechnology, epidemiology, public finances, labor market, etc.

propensity-score weighting (like in matching regression in Coarsened Exact Matching method or IPW in EWM method); and yet another time, as an input to pre-estimating the sample individual *average treatment effect on the treated* (ATET). This latter procedure allows for different pre-estimation of welfare gains impact for each sample observation *depending on whether* the observation belongs to the ex-ante treated (policy-assigned) portion of the sample population or to the non-treated counterpart. This technique encapsulates a group of estimators whose statistical properties have been well-documented elsewhere in a branch of theoretical statistics literature called *projection and imputation* techniques (Abadie and Cattaneo, 2018 in particular Chapter 4.6; Lechner, 2023 with empirical extrapolation; Wager and Athey, 2018 and Athey and Imbens, 2017 for earlier unifying reviews).

The use of doubly robust estimator bestows upon Policy Learning, on the one hand, superior statistical consistency where in particular the corroborative effects of data imperfection (such as missing data in tabular data types very frequently encountered in public finance and labor economic empirical applications; my implementations in Chapter 5.3 serving as illustrative examples) are given alleviating layer of safety because as long as either one channel of the twofold dataset-leveraging technique described above remains sufficient then the whole estimation remains statistically consistent—a hindrance other methods such as IPW would have disqualified EWM from using the same imperfect dataset. As before in EWM, mathematical proof is here of secondary aim to the chapter and thus referred to either Chernozhukov et al. (2018), Athey and Wager (2021), or Knaus (2022), all of which offer excellent overview to the use of doubly robust estimator in statistical decision theory for optimal policy assignment. On the complete development history of the statistical literature leading to the doubly robust estimator as representative of the group of *hybrid* methods combining propensity-score weighting and projection/imputation tools, I emphasize the excellent review by Abadie and Cattaneo (2018).

Chapter 4.4 Methodological Extensions

Methodological Extension I: Other ML Algorithms

As mentioned above, the majority of empirical implementations of the methods for treatment assignment using statistical decision theory has opted for the statistical learning algorithm called Random Forest to execute its doubly robust estimation steps. As Athey and Wager (2021) most explicitly hints upon in the case of Policy Learning, nevertheless, the asymptotical theoretical foundations do not preclude the use of more recent, potentially more powerful machine learning algorithms such as Neural Network (NN) or Support Vector Machine (SVM).

Comparably, the summary on the use of doubly robust estimation techniques for improving treatment assignments provided by Knaus (2022) includes a list of potential other algorithms they consider as hypothetical replacement for the default algorithm in the model. More specifically, calling the general method *DML_aipw* (technically: double machine learning, augmented inverse probability weighting), the study paper at least four further machine learning algorithms that it deems suitable replacements: OLS (i.e., standard ordinary least square regression), Ridge (a variant of penalized regression), LASSO (another variant of penalized regression with additional shrinking/variable-reduction feature), and generalized random forest (GRF, a generalized extension of the decision trees).

In my own empirical implementation in Chapter 5, I illustrate the use of Neural Network (also known as Deep Learning) as replacement for Random Forest in the Step 1 of Policy Learning method for the analysis of the job market SIAB dataset. To the best of my knowledge, I present thereby the first empirical implementation that incorporates such methodological extension and reports its resulting advantages.

Methodological Extension II: Interpretability Methods

While powerful in their predictive power, in the past machine learning techniques had been adopted into the literature of economics at a more careful speed compared to the rate at which they are increasingly ubiquitous embedded in, to name an example, medical and bio-engineering literatures. While recently the adoption of advances from the literature on statistical learning into econometric models has accelerated (most comprehensive outlook on this phenomenon is given in Athey and Imbens, 2019), the prudence with which economists have proceeded in incorporating especially the so-called *black-box machine learning models* is well warranted when one takes into account the legal, political, and ethical realities the economic literature as a scientific field must successfully manage to navigate—as much as, if not arguably mostly more so, than in other disciplines. This rationale is particularly palpable in the public economics literature and its intersections with adjacent fields concerning direct program interventions to individuals as immediate beneficiary groups.

A machine learning model is referred to as having 'black box' characteristic when it does not reveal its internal mechanism to a sufficient degree of simultaneous transparency, i.e., its end results (prediction accuracy, as in computer science literatures; or policy recommendation on treatment allocation, as depicted in this and the following chapters) are not immediately traceable back to any specific step or parameter input in its preceding implementation process. On the other hand, several machine learning models are *inherently interpretable* – chief examples among this category includes decision trees with depth 1 (an example of which is given in this dissertation in Chapters 5.3 and 5.4) and linear regressions, including its extensions with either propensity-score weighting or matching methods (for a recent example of application in the intersection between literatures on public finance and tertiary education, see Herberholz and Wigger, 2021).

A very recent addition to the machine learning literature is a growing body of research on *interpretable machine learning* (IML), which seeks to supplement black-box machine learning methods with modifications that can help explain the the technical process with which an algorithm reaches its conclusion. These interpretability modifications are typically applied to the machine learning methods *after* the sample data had been used in the *training* step, which essentially means that the IML modifications can*not* interfere with the machine learning method's initial interaction with the provided dataset; such a contamination could have otherwise hindered the efficiency and even applicability of the method overall. This aspect is crucial especially in empirical implementation settings such as those studied in Chapter 5 of this dissertation where the datasets are, though remarkably large, official administrative records—unlike many instances in computer science studies, where datasets such as e-commerce sales data or online user interactions can be replicated or (re-)simulated abundantly.

With this background in mind, I concluded that incorporating interpretability modifications constitute indispensable avenues for empirical studies seeking to implement statistical treatment assignment methods such as Policy Learning and Empirical Welfare Maximization. From the whole array of currently available IML methodology summarized most recently in Molnar (2022), two interpretability techniques lend themselves most relevantly to the four methods of statistical decision theory evaluated in Chapter 4.3. First, the *variable importance* interpretability method can complement either Random Forest or Neural Network, which both belong to a group of machine learning models from which I can choose to proceed the treatment assignment framework[29]. The goal of this technique it to de-parallelize the permutations of the interactions between features (i.e., the candidates for independent variables), which the machine learning algorithm executes 'in the background' (in data science, the term refers to components of algorithm steps that are calculated automatically by the machine

[29] Recall from preceding subchapters that Random Forest, Boosting, and Neural Networks are three multi-layered machine learning methods that can be interchangeably chosen for executing Step 3 of Policy Learning.

learning method but are not displayed) and utilizes to achieve its prediction, and to extract from these feature interactions information on which variables have affected the others most frequently while the learning algorithm was performing its task "in the background'. Using these numbers, the interpretability method then constructs a list of either 5 or 10 *most important variables*, i.e., the most impactful variables. This shortlist of ranked variable importance can then be exploited in a supplemental analysis complementing the main treatment assignment framework, such as to be input elements for heterogenous treatment effects (HTE) or subgroup analysis. Indeed, recent comprehensive empirical implementation of statistical decision theory for treatment assignment such as Knaus (2022) shows how performing additional heterogeneous treatment effects (HTE) analysis complements the policy interpretation of Policy Learning. Knaus (2022), however, incorporated no interpretability modification to improve the treatment assignment problem. In Subchapter 5.3.1, I show how the *variability importance* can be supplemented directly into the Policy Learning method.

The second interpretability modification that can be supplemented to a treatment assignment method is the *accumulated local effects plot* (ALE plot). This method has the advantage of calculating differences in predictions instead of their averages, thereby preventing effects of other independent variables from contaminating the estimation of its impact—a common issue discussed in the literature of machine learning (see, e.g., Bohren and Hauser, 2021). I illustrate the use of ALE plots in Subchapter 5.3.2.

Chapter 4.5 Concluding Remarks

This chapter summarized recent methodological advances in the growing literature of treatment choice, at the intersection of current literatures in economics, econometrics, and data science. It highlights the advantages of utilizing insights from statistical decision theory into the setting of program (re-) distribution, delineated recently by the statistical methods Empirical Welfare Maximization (EWM), Penalized Welfare Maximization (PWM), Policy Learning and Asymptotic Minimax Regret (AMMR). Particular emphasis is given on the methods EWM and Policy Learning, which together form the methodological basis for the empirical implementations in Chapter 5.

This chapter also shed light onto several key areas of methodological development in the literature, namely the extension of the doubly robust estimation at the heart of aforementioned statistical-decision-theoretic models with newer machine learning algorithms such as Support Vector Machine (SVM) and, in particular, Neural Network (NN) a.k.a. Deep Learning. Particularly important for implementations in public finance literature with regards to public program assignments, Chapter 4.4 also discussed the insightful development in the small but growing literature strand of interpretable machine learning (IML). In each case of methodological extension, a short preview was given to the particular aspects of development which will be implemented directly in empirical settings in Chapter 5.

Fruitful future research avenues include the potential incorporation into the statistical treatment allocation methods the so-called *local agnostic interpretability* modifications such as Shapley value, which would complement the use of variable importance and ALE plot discussed here. Moreover, multitude of empirical applications of the methods and would potentially lend further credibility to the methodological applicability. The next chapter represents my own efforts toward contributing to this direction.

Chapter 5

Empirical Implementations of Policy Learning

Comparative Study: Empirical Implementations of the Method and Its Extensions (from Chapter 4) on Two Administrative Datasets

Chapter 5.1 Introduction

While the upsurge of data availability represents excellent research material, public economics needs adequate update to its array of analytical tools to meet the massive increase of empirical complexity. In the last two decades, the fields of statistics and data science have produced remarkable innovations in data analysis at remarkable speed. As a result, there is a real and huge potential for incorporating these progresses into existing methods of public economics. To be specific, advances in machine learning that have previously revolutionized medical sciences by using large individual records to predict treatments and help improve practitioners' accuracy will be able to assist policymakers in choosing optimal treatment allocation under individual heterogeneity. The advantages to this approach are at least threefold. First, the sheer size and technical complexity of new administrative datasets render conventional methods unsuitable for determining optimal recipient criteria and beneficiary groups. For example, the Policy Learning method developed by Athey and Wager (2021) and touched upon in Chapter 4 uses random forests to immediately learn and predict two straightforward optimal policy criteria based on observational data with 28 individual covariates and over 19,000 observations. If done manually using existing methods, this task would have required weeks if not months of costly trial and errors before arriving at the desired criteria combination. More demanding still from computational point of view is the fact that the final optimization algorithm is essentially a non-convex optimization problem[30].

[30] Predecessor methods such as Empirical Welfare Maximization discussed in Chapter 4 as methodological benchmark to Policy Learning *only enable* binary, linear decision rules as policy

The second benefit appeals directly to issues hugely relevant to the public economics literature, namely the unwavering demand for transparency on the one side and political or ethical constraints on the other side. Using machine learning to produce optimal policy and its distribution criteria satisfies both, because sensitive individual attributes such as gender or race can be fully excluded from the final policy recommendation despite being incorporated into the algorithm in the previous stages. Recall, however, that the very possibility of leveraging the ever-widening range of available variables in the public administrative data constituted the raison d'être of statistical decision theory – hence the inclusion of these sensitive variables in the early stages (e.g., Step 1–2 in Policy Learning). More fundamentally, because it has been empirically observed that in many public program settings these variables "seem to be confounders" (cf. Knaus, 2022; Athey and Wager, 2021), without whose inclusion the Conditional Independence Assumption[31] that underpins the causal credibility of the statistical decision theoretic methods assumption will be compromised. In other words, Policy Learning allows us to alleviate the confounding problems when early descriptive analysis dataset reveals an imbalance in the dataset pertaining to segregations that would have been legally/ethically ineligible had they been used as policy divisions. For example, because their heterogeneity analysis revealed a race-based asymmetry in the distribution of job training in the so-called Greater Avenues for Independence (GAIN) program in the United States between 1986 and 1993, Athey and Wager (2021) *included* race in their de-confounding step (Policy Learning Step 1) but *excluded* them from the pool of considered decision variables in the final, optimization step (Policy Learning Step 3).

---

class, which can be constrained into convex optimization problems and thus be solved using older, non-machine-learning-supported optimization algorithms. Allowing for even simple alternative rules such as (depth-1) decision trees, however, would already enlarge the optimization problem into *non-convex* territory and render the EWM infeasible—let alone using more sophisticated alternatives such as Random Forest or Deep Learning. Yet these latter alternatives are crucial if one were to leverage the full potential of the Big-Data-sized datasets—hence the rationale for using Policy Learning and its further modifications such as the ones I developed in Chapter 4.4 and utilized in the next subchapters.

[31] Also known in the literature as other aliases, e.g. as causal identification strategy in causal regression and/or matching methods: unconfoundedness, exogeneity, selection-on-observables, ignorability.

Finally, and crucially, whereas methodology of inferential statistics often emphasizes the need for field experiments in order to achieve minimum amount and quality of data, in reality these experiments are in most cases either too costly to perform or infeasible due to political or ethical constraints, or both. Machine learning methods, on the other hand, are sufficiently equipped to be able to learn optimal policy even in the absence of experimental data. In other words, the aforementioned, newly available observational data in various administrative forms can be readily utilized to learn and predict optimal future policies.

Building on the literature strand pioneered by Manski (2004) and summarized theoretically in Chapter 4 on the use of statistical decision theory for treatment choice problems in econometrics, this Chapter 5 implements statistical treatment decision theory in the setting of two administrative datasets. First, Germany's federal-administered vocational training program is modeled as statistical decision problem using the Empirical Welfare Maximization and Policy Learning methods. Secondly, similar treatment is given to the newly available dataset of Indonesian health insurance. For the latter, methodological modifications were undertaken and are reported here, pertaining to the dynamic assignment problem of the multi-level treatment of the healthcare insurance setting.

In doing so, I synthesize latest methodological advances in statistical learning to obtain and evaluate statistical treatment rules that maximize welfare while respecting exogenous (political, ethical) constraints. Aggregate welfare gains expected from adopting these rules instead of manual allocation are reported.

The remainder of this Chapter is organized as follows. Subchapter 5.2 briefly sheds light on the theoretical underpinnings that microfound the implementation depicted in Subchapters 5.3 and 5.4, while avoiding unnecessary redundancy with Chapter 4 by highlighting the key methodological elements where judgement calls on certain parameters are needed. While my technical choices on these parameters set as input into the algorithm for the empirical implementations are specified in the later subchapters, in general there are lessons to be concluded for future implementation, which are then succinctly reiterated in the remainder of the Chapter.

Chapter 5.2 Methodological Background

This subchapter enlists the core methodological aspects that are utilized in the implementations in Subchapters 5.3 and 5.4. I highlight key methodological developments along the recent literature strand which adopts statistical decision theory into treatment assignment problems. A majority of them study sample data generated through randomized experiments, though a recent few manage to leverage observational data as well. Seen from the practical perspective of policymaking in adult education landscape, the latter fits particularly well with the configuration made necessary by many real-world settings of administrative labor market data. I begin with the central ideas underlying statistical treatment rules.

Following Manski (2004) and the literature it subsequently inspired, the core problem in treatment assignment from the perspective of statistical decision theory is one where a policymaker must decide on some configuration of treatment allocation for the population while contending with incompletely observed heterogeneous treatment effects and maximizing expected average welfare. Her approach is necessarily of second-best nature, since the first-best solution is unattainable due to incomplete information over the true distribution of treatment response.

To proceed with her task, the policymaker can leverage information gained from a sample data of the population. Recall the definition of statistical decision functions introduced earlier as a function that maps each individual sample treatment response to a decision, with statistical treatment rules describing such functions in the specific context of treatment assignment. Regret is then defined as the difference between the (unachievable) first-best outcome and the outcome yielded by the best-in-class policy generated by the chosen statistical treatment rule. In this regard, regret in the sense of statistical decision function is akin to loss functions in prediction methods, such as the classic mean-squared error. Again, the important advantage of statistical treatment choice is that it incorporates the welfare weighting directly into the optimization step, as opposed inferring parameters first and then plugging them in at a separate stage. In doing so, certain exogenous constraints faced by the policymaker can be integrated – conversely, the aforementioned approach of individualized treatment effects in the machine learning literature cannot allow such constraints.

The principle, based on which the procedure learns optimal treatment assignment through generalizing sample-based information as if it applies analogously to the population, is called empirical success rule (Manski, 2004). Kitagawa and Tetenov (2018) extend this concept to allow for constraints on the set of feasible treatment policies, such as when potentially discriminatory attributes (gender, migration background, etc.) are required to be excluded from the pool of covariates upon which the algorithm may learn optimal treatment rules. In reference to the seminal work on empirical risk minimization (ERM) in early machine learning literature (see, e.g., the summarizing contribution by Vapnik, 2000), they call their statistical treatment allocation method empirical welfare maximization (EWM). Mbakop and Tabord-Meehan (2021) develop the EWM method further by introducing penalized welfare maximization (PWM) method, which expands the eligible forms of policy constraints as well as introduces an alternative approach for inherent model selection using the hold-out procedure.

In many settings with administrative data and for vocational training programs in particular, as several studies aforementioned in the literature review pointed out, there is a palpable need for methods that can leverage not only data from randomized controlled trials but also observational data. For example, Germany's largest state-administered labor market dataset *Sample of Integrated Labor Biographies* (SIAB) documented over 66 million individual observations with 31 attributes including recorded participation in vocational and professional training programs between a timespan of over four decades. The dataset offers plentiful research opportunity regarding individual heterogeneity yet does not explicitly address any particular experiment. Instead, the researcher is called to seek a way to exploit the rich observational data and incorporate it into the statistical decision theoretic framework. As briefly mentioned before, Athey and Wager (2021) expand the statistical treatment rules in this exact direction: their policy learning method uses matching procedure with doubly robust scores to compensate for the lack of counterfactual treatment response and delivers optimal treatment allocation based on the resulting propensity score. Similarly to the empirical maximization method, though, their policy learning method also allows for exogenous policy constraints.

Chapter 5.3 Empirical Implementation I: SIAB

This subchapter depicts the first of two empirical applications of the Policy Learning framework for two respective empirical settings of treatment allocation[32]: first, with regards to active labor market policy (ALMP, as the abbreviation is commonly known in the literature), and secondly, in the context of state-administered health insurance. The latter application benefits from the relatively novel dataset on BPJS-KIS, the Indonesian public health insurance, for the research use of which I successfully attained official permit in 2021. The empirical application in active labor market policy setting, on the other hand, leverages the dataset SIAB of Germany's Federal Institute for Employment Research (*Institut für Arbeitsmarkt- und Berufsforschung*, IAB). Subchapter 5.3.1 first delivers an overview of the current state of the literature –primarily of the literatures– on the active labor market policies in Germany, offering relevant institutional background and pointing to recent literatures in public finance and labor economics which are most relevant. Section 5.3.2 then recounts the empirical implementation of the Policy Learning methodology on the aforementioned SIAB-dataset—without redundantly repeating the methodological explanations given already in Chapter 4; rather highlighting several key technical decisions that I tailor-suited with regards to the pre-requisites of this particular dataset. To the best of my knowledge, this dissertation contributes to the literature strand the first result of statistical treatment allocation applied for the German labor market setting.

[32] As clarified in the *Subchapter 5.2* and more extensively in *Chapter 4*, the terminology of *treatment allocation* does not limit itself to applications in health economic setting, where the term *treatment* could refer to another, narrower-sense usage. Instead, as is equally common practice in the econometric literature of at least the last three decades (e.g., in literature of randomized control trials/RCTs), treatment allocation in the literature of statistical decision theory refers broadly to *policy assignment* the state planner/government makes in various public sectors.

Chapter 5.3.1 Institutional Background and Dataset

A comprehensive, relatively recent summary by Klös (2021) provides overview of Germany's current landscape of vocational and professional education as well as critical review of related policies, both those already implemented and others subject to current debate. Drawing comparison with the state's growing proportion of financial support in research and higher education, Klös argues that the pivotal role of craftmanship and skilled manual labor for German small to medium-sized enterprises justifies a parallel increase in public spending on job trainings. This notion reflects a perceivable consensus among labor and education economists over the need to extend the scope and quality of adult education policies by introducing innovative programs while instilling more rigorous policy evaluation methods (see., e.g., *Data Report on Vocational Education* by the Federal Institute for Vocational Education (Bundesinstitut für Berufsbildung, 2020)).

Though researchers share a rather common view on the importance and necessity to expand the scope and strengthen the impact of public vocational programs, they do not always agree on which policy levers to pull. Card, Kluve, and Weber (2010, 2018) provide comprehensive review of adult education policies across several countries with focus on how they affect labor market outcomes. They point out the presence of heterogeneity across different participant groups in previous program evaluations in the literature – many of which remain unobserved in the original studies, thereby hampering the precision of the initial estimates of the respective program's effectiveness towards its intended labor market goal. Particularly, findings from several studies suggest minimal or non-existent causal effects of specific program types on labor market outcome. For example, Görlitz and Tamm (2015) documented no significant effects of German short-term training vouchers introduced in 2008 on labor market outcomes such as employment rate or wage[33]. Evaluating similar program in Switzerland, Schwerdt, Messer, Woessman, and Wolter (2012) 's caution against misleading outtake on efficiency when

[33] Görlitz and Tamm (2015) did find a weak positive impact of the program on perceived skill match, a survey item which was gauged through self-reported subjective assessments of participants. The extent to which this effect can convincingly be used to inform future policy design remains unclear.

heterogeneity is not sufficiently considered in both treatment effects as well as disparity between intention to treat and actual intake rate across all potential beneficiaries.

To be sure, the problem of treatment effects varying in their magnitude or even in some cases diverging in their direction as a result of residual heterogeneity across individual characteristics is not unique to the active labor market policy or adult education literatures. More generally, along with the advent of increasingly large and deep (displaying observations often down to individual features) administrative data, researchers across economic and social-science disciplines have come to acknowledge the need to address heterogeneity in both population and sample datasets that cannot be captured through conventional identification methods. In some cases, accounting for heterogeneity can open up new line of evidence where the literature had previously followed an opposite consensus. For an example from the labor market sector, Caliendo, Hujer, and Thomsen (2008) show that while prior studies found negative average treatment effect for German job creation schemes (JCS), disaggregating enabled them to present a clearer picture of the situation: While their results affirm negative effects for most treatment groups, there was clear indication of positive effect for a particular beneficiary group, namely the long-term unemployed. A subsequent policy recommendation could benefit from this extra information by targeting the narrower recipient group – an advantage that would have been foregone had the literature not adopted heterogeneity tools.

Finally, I note that impetus towards identifying optimal treatment assignment rule for a given policy can be traced back among several previous studies in labor-related settings. Huber, Lechner, Wunsch, and Walter (2011) evaluated multiple German active labor market policies and highlighted heterogeneity as hidden problem. Lechner, Miquel, and Wunsch (2011) exploited individual heterogeneity in rich administrative data to differentiate their analyses for of short, middle, and long-term effects of various German adult education programs. These and other existing studies commonly present their account on heterogeneity in auxiliary section on sensitivity analysis, as a means to strengthen the study's main claim (for example, Busse et al. (2017) explicitly describe the selection procedure of their eligibility rules as manual selection). The choice of alternative assignment rules thus in the first place, despite accompanied with adequate justification, were often less than perfectly systematic. On the contrary, statistical treatment rules would allow the policymaker to reduce time and effort otherwise needed in building and evaluating alternative scenarios.

I utilize a sample version called SIAB CF of Germany's administrative dataset *Sample of Integrated Labor Biographies* (SIAB), which as a 2% sample of the IEB registry contains over 66 million individual observations with 31 covariates between 1975–2019. SIAB facilitates observation on the impact of further training (*Förderung der beruflichen Weiterbildung)* program on labor market outcomes, of which I use earnings in the post-treatment periods as dependent variable. This program is distinct from other types of Germany's (also state-administered) active labor market programs, such as retraining (*Umschulung*) programs or integrations subsidies (*Einarbeitungszuschuss*) (for more detailed look on this categorization see, e.g., Fitzenberger & Speckesser, 2007).

Kitagawa and Tetenov (2018) and Athey and Wager (2021) each illustrated their proposed method on a historical active labor market policy based on randomized experiment: California's Greater Avenue for Independence (GAIN) welfare-to-work program (RCT in 1988–1993) in the case of Athey and Wager (2021) and the US national Job Training Partnership Act (JTPA) experiment in the case of (Kitagawa and Tetenov, 2018). Under these premises, the propensity score (the likelihood of being treated given certain covariate values) is known to the researcher and the resulting treatment allocation rules inherit the causal nature of the estimated CATE under the *unconfoundedness assumption*.

The *Förderung der beruflichen Weiterbildung* vocational policy in the SIAB dataset that I model as the treatment, on the other hand, contains neither a natural nor field experiment. Athey and Wager (2021) established theoretical guarantee that asymptotic optimality could still be achieved in settings with pure observational data, provided one uses recent advances in the statistical learning literature called double/de-biased machine learning (where the propensity scores are now also estimated from the data). Empirical implementation of this remains scarce, however (an exception is Knaus, 2022, who re-examined non-experimental Swiss job training program also using doubly robust estimation). To the best of my knowledge, this dissertation thus represents this literature strand's first result for the German labor market.

With regards to the Scientific Use File (SUF) of the dataset Sample of Integrated Labour Market Biographies (SIAB), it represents a dataset technically classified as the factually anonymized[34] version of the 2% random sample drawn from the Integrated Employment Biographies (IEB), which covers an extensive time period between 1975–2017 and comprises in total over 1.8 million individuals in German who fulfills at least one of the following employment status: employment subject to social security (in the data since 1975), marginal part-time employment (in the data since 1999), benefit receipt according to the German Social Code III or II (SGB III since 1975, SGB II since 2005), officially registered as job-seeking at the German Federal Employment Agency or (planned) participation in programs of active labour market policies (for further description see, e.g, Antoni et al., 2019).[35]

The dataset SIAB combines a massive number of recorded observations with an extensive coverage of individual characteristics covariates including for example the exact time period of unemployment and subsequent re-entry to active labor market, gender, last attained education level, as well as average income). As a result, this material is highly suitable for an implementation of the Empirical Welfare Maximization method and Athey and Wager's (2021) policy learning

[34] Technical definition and the ensuing data management process of factual anonymization, particularly in the manner in which the research data center (*Forschungsdatenzentrum*, FDZ) of the IAB executes this, can be read in FDZ's Datenreport.

[35] [35] While my research project was granted official access to SIAB-R 7519, an anonymized sample version of the SIAB source dataset, computational constraints along the way have proven it necessary to limit my implementation to a smaller sample version of SIAB designated SIAB CF. This dataset recounts in turn 99,284 unique observations for 6,741 individuals across 21 variables.

algorithm. In my work, I used the available covariates as input variable, from which the random forest can predict the optimal decision trees for allocating publicly administered job training. The average daily income was excluded from this pool of covariates, since it is used as the dependent variable to be maximized. As previously mentioned, several politically or ethically sensitive variables can still be used during the machine learning procedure, despite being excluded in the last optimization stage in order to prevent the inclusion of any sensitive criterion in the resulting decision trees.

Chapter 5.3.2 Implementation on SIAB Dataset

Results: Optimal Policy Assignment (in Decision Trees)

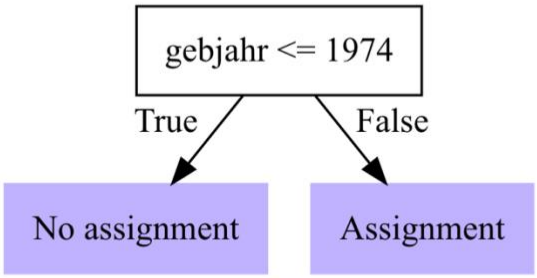

Figure 5.1 *Optimal decision tree obtained through Policy Learning. Depth = 1.*

Corresponding to the machine learning model technique selected[36] for benchmark implementation for the (re-)assignment of active labor market program in the contextual form of job training, Figure 5.1 depicts the result of the implementation of Policy Learning involving all Steps 1–3 yet without blending out sensitive variables of *age* and *gender* in the final step. This resulting policy recommendation in the form of a depth-one decision tree one (with one binary decision node) can be interpreted directly as to prescribe the policymaker to split his decision on (re-)administering the job training allocation based on one single decision question—whether the evaluated individual was born later or before the year 1974 (*gebjahr* stands for year of birth). According to this simplest decision prescription, the aggregated (utilitarian) social welfare would have been maximized were the job training be assigned to individuals within the sample who are not older than born in 1974. This example exemplifies a case where the Policy Learning method indeed would recommend a re-assignment of the job training program based on year of

[36] Recall that empirical implementations in the literature so far of statistical treatment allocation methods Policy Learning, EWM, PWM and AMMR conservatively stick to either singular decision trees or random forest, yet the methodological framework does *not* preclude incorporating even more sophisticated machine learning methods such as neural network (NN; also known as Deep Learning) and support vector machines (SVMs), the latter of which are conceptually closest to SDT-esque categorical prediction tasks yet still inadequately adjusted for Policy Learning. Chapter 4 discusses all these considerations..

birth (i.e., age), which could be perceived as discriminatory and is thus rather legally/politically unfeasible.

As emphasized conceptually before, one key advantage of the Policy Learning method is that it facilitates the policymaker to specify in Step 3 which sensitive variables the algorithm should leave out from the pool of possible decision nodes (again, without having to sacrifice their utilization in the prediction-related Steps 1–2). Intuitively, the algorithm steps resulting in this 1-depth decision tree could be further modified with regards of the complexity of the recommendation. The aforementioned tree depths can be increased by the policymaker in Step 3 of the Policy Learning algorithm into 2 (in the most groundbreaking studies recently even into depth = 3; though the very nature of the complexity of both the graphical portrayal *and* the interpretation for policy becomes practically counterproductive especially in settings of public finance, where strands of literature have been dedicated to unravel the adverse behavioral response effects of complexity in regulatory or taxation policies). With a depth of 2, however, the policy recommendations obtained by the Policy Learning method achieves a justifiable balance between clarity and the prediction power attained by allowing for more complicated statistical algorithms. Moreover, as exemplified by the results seen in Figure 5.2, the moderate complexity of this configuration renders itself especially comparable to conventional, manually-picked decision criteria for administering such public ALMP programs without statistical decision theory (see, e.g., Fitzenberger and Speckesser, 2007) and Lechner et al., 2011).

The resulting decision tree is presented in Figure 5.2, with neither of the sensitive variables *age* nor *gender* picked as decision node and instead two other decision variables are used: *persnr*, which codifies the identification number administratively assigned by IAB (and thus already undergoing the anonymization process); and *tage_svb_gr*, which recounts the duration the person was in active labor employment (number of days as measurement unit). The decision tree can thus be interpreted as re-prioritizing the training program for two subgroups of the sample population: 1) individuals with *persnr* between 2719

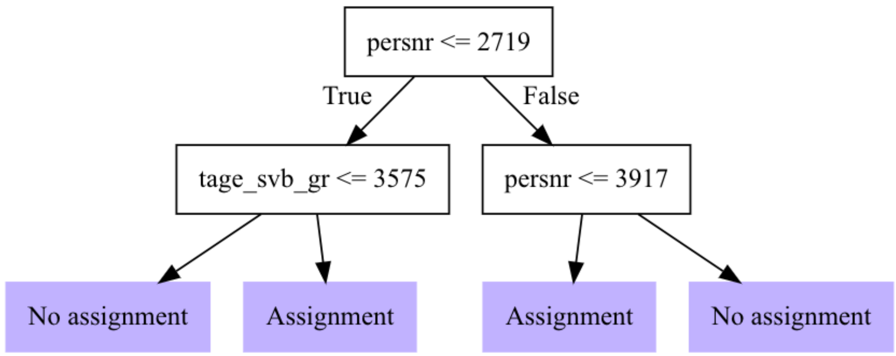

Figure 5.2 *Optimal decision tree obtained through Policy Learning. Age and gender excluded from Step 3. Depth = 2.*

persnr = identification number (administratively assigned)
tage_svb_gr = number of days in active labor employment during the last 10 years

and 3917[37] and 2) individuals whose *persnr* smaller than 2719 *and* who were registered as unemployed for less than 2 months in the past 10 years (3575 days ≈ 9 years and 10 months in employment).

Despite the disadvantages I pointed out for the results with tree depths of 3 mentioned before (complexity, impracticality of enacting), as a conceptual inquiry and as a sensitivity analysis to my preferred specification leading to Figure 5.2 (tree depth = 2 | decision nodes = 3 | sensitive characteristics blended out), a policy prescription with the specification (tree depth = 3 | decision nodes = 3 | legally-constrained characteristics *not* blended out due to sample constraint when exploited for depth-tree of 3) is reported below.

[37] The policymaker desiring deanonymized individual identification could then, in a separate step not belonging to the scope and aim of this dissertation, potentially reach out to IAB to deanonymize the *persnr* numbers by linking to a source dataset—a step of highest data security requirements that, to the best of my knowledge, IAB grants only for government's purposes.

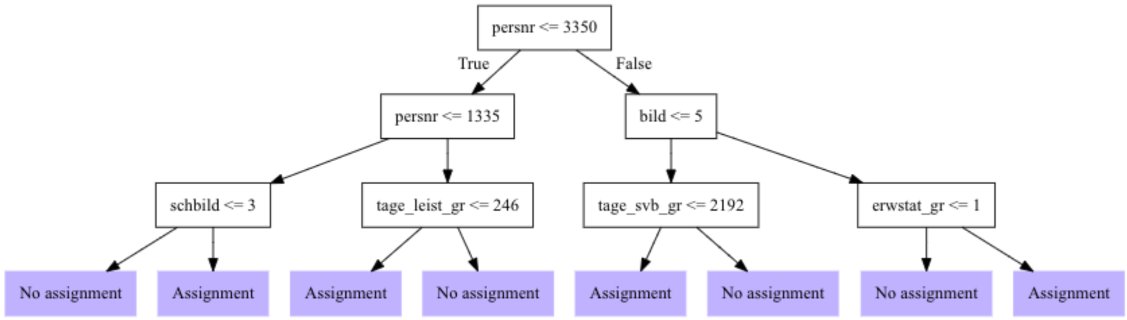

Figure 5.3 *Optimal decision tree obtained through Policy Learning. Age and gender excluded from Step 3. Depth = 3.*

persnr = identification number (administratively assigned)
tage_leist_gr = number of days as recipient of social transfer recipient in last 10 years
tage_svb_gr = number of days in active labor employment during the preceding 10 years
erwstat_gr = current employment status
schbild = years of primary and secondary education
bild = years of post-secondary education

As Subchapter 5.5 will discuss, it turns out that in the SIAB-CF setting, the average gain yielded by depth-3 was in fact virtually equivalent to the one yielded by its simpler, depth-1 counterpart. From the point of view of policy recommendation, therefore, the use of depth-1 and depth-2 Policy Learning turn out to be preferrable in this empirical setting.

Chapter 5.4 Empirical Implementation II: BPJS-KIS

Chapter 5.4.1 Institutional Background and Dataset

I analyze an observational dataset of Indonesian health insurance program *BPJS Kesehatan*. Enacted in January 2014, it represents the first state-administered health insurance in the country's history. As of 2019, 224 million Indonesians were enrolled in the program (83% of the country's population). As a developing country, Indonesia is currently developing its social insurance structure in an accelerating speed, assisted by the advent of barrier-breaking technologies such as big data and machine learning. For the first time since the country's independence in 1945, a national health insurance system was introduced in 2014, followed by pension and unemployment insurances. In 2019, Indonesian authority announced the availability of large-scale dataset upon application for suitable research purposes. While this sample dataset contained individual health complaints and medical treatments 2015–2016, an updated version with significant extensions in the level of details and coverage was released very recently in December 2020. The latter covers an extended time period from 2015–2018 and covers multiple stages of medical treatments for over 1.9 million individuals, 9 million observations and 47 covariates (of which 15 are of administrative labelling nature, thus around thirty covariates were filtered in the manual pre-processing before being input into Policy Learning Step 1).

I characterize the two distinct referral schemes "FKRTL" (first treatment stage; can be translated from the Indonesian most closely as "with referrals") and "FKTP" (second treatment stage; "non-capped") as a multi-treatment allocation problem (for definition and discussion of the statistical properties of such problems, see, e.g., Zhou et al., 2018). In total, 4.317.826 medical events (non-capped) and 1.598.642 (with referrals) were reported, constituting the multi-stage treatment assignment scheme. The objective function, i.e., the welfare gains later reported in Table 5.2, is to be interpreted as the efficiency gains (in the equivalent unit referred to in the variables representing FKRTL and FKTP) from re-prioritizing the type of treatment (two-stages, with both FKTP and FKRTL present *or* one-stage, *directly* to the referrals to advance specialist treatment centers—the FKRTL).

The policy recommendations in forms of decision trees are reported below, with the interpretation approach analogous to those introduced in Subchapter 5.4.1. Table 5.2 in turn report the aggregated total sum (from either FKRTL or both levers of FKTP and FKRTL) of achievable gains obtained through re-arranging the treatment assignments by leveraging the sample population's individual variables by means of Policy Learning in its differentiated configurations involving methodological extensions elaborated in Chapter 4.4. Specifically, the implementation of Policy Learning (and for that matter any of the portrayed benchmark statistical decision-theoretic methods) for the health insurance setting of BPJS-KIS necessitates a modification to facilitate *multi-action* or *multi-(assignable)-treatment*.

To reflect the multi-stage treatment assignment mechanism observed in the BPJS setting, I model the assignable treatments as follows. For better comparisons with preceding, descriptive national studies (e.g., Ariawan et al., 2020; Chendra, 2020), the original Indonesian variable names as used in the dataset are listed as well.

Treatment name	Dataset variable	Description
Capitation-treatment only	FKTP Kapitasi	Low-cost treatments at first health institution. Capped at insurance-defined maximum costs. No further referrals.
Non-capitation treatment	FKTP Non-Kapitasi	Advanced treatments directly at first-visited institution, costs beyond insurance cap. No further referrals.
Referral treatment	FKRTL	First-visited health institution issues referral to a second (specialized) institution. Indicates multi-stage treatment. Costs beyond insurance cap.

Table 5.1 Assignable treatments and their corresponding categorization in my Policy Learning implementation on the BPJS health insurance dataset.

Chapter 5.4.2 Implementation on BPJS-KIS Dataset

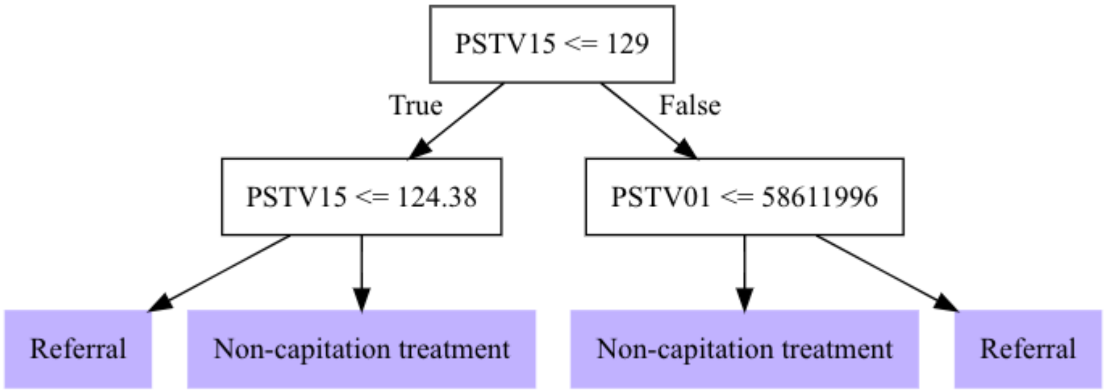

Figure 5.4 *Optimal decision tree obtained through Policy Learning using Random Forest configuration for BPJS health insurance setting. Depth = 2*

Corresponding to the machine learning model technique selected for benchmark implementation for the (re-)assignment of active labor market program in the contextual form of job training, Figure 5.4 depicts the result of the implementation of Policy Learning involving all Steps 1–3 yet for the multi-program setting of Indonesia's health insurance dataset BPJS-KIS, with the designation of the outcome of interest and possible treatments discussed in Subchapter 5.4.1.

The implementation of Empirical Welfare Maximization without blending out age and gender variables (i.e., as in the original method of Kitagawa and Tetenov, 2018) as benchmark yields a recommended policy tree as depicted in Figure 5.5.

Finally, Figure 5.6 reports a depth-2 decision tree result with the Step 1 of the Policy Learning algorithm modified by using neural network in the estimation step of the prediction parameters, which the literature strand calls nuisance parameters, consisting of two predicted parameters namely an estimated propensity score (likelihood of a given individual to be treated) and an estimated potential outcome (ex-ante/before the re-allocation algorithm which take place in Step 3). In other words, I substitute the use of random forest as default prediction method in the *policytree* package (by Athey and Wager, 2021 and introduced in Chapter 4) with a neural network prediction, executed using in the R statistical software using *keras* package. In effect, I present here first empirical evidence of viability of using neural network in Step 1 of Policy Learning – the theoretical notion of which has, as aforementioned, been hinted upon as further research avenue by inter alia Athey and Wage (2021) themselves.

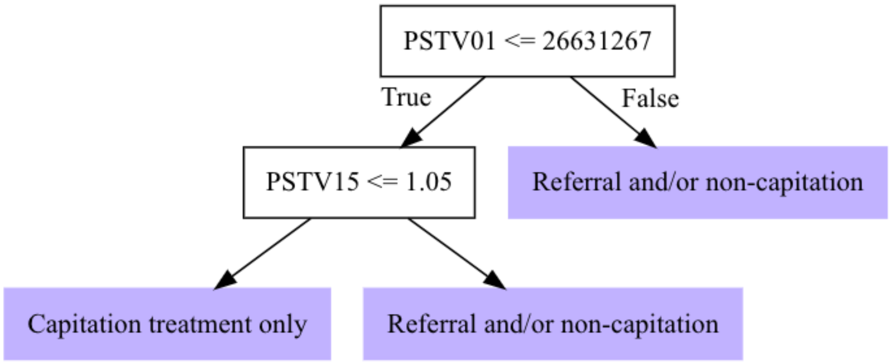

Figure 5.5 *Optimal decision tree obtained through EWM (as benchmark method) for BPJS health insurance setting. Depth = 2*

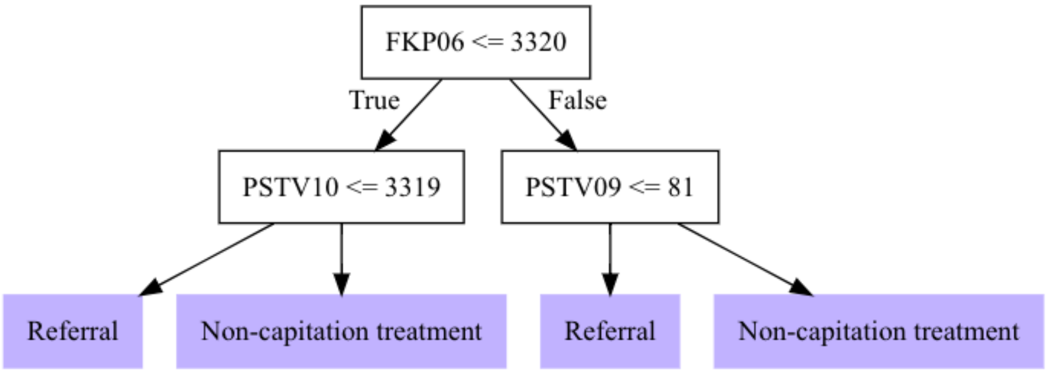

Figure 5.6 *Optimal decision tree obtained through Policy Learning using Neural Network (as methodological extension) for BPJS health insurance setting. Depth = 2*

FKP06 = city of residence, as manually re-registered at first health institution visit
PSTV10 = city of residence, coded administratively at health insurance take-up
PSTV09 = province of residence

Chapter 5.5 Results

As the key advantage from using the Policy Learning and Empirical Welfare Maximization as statistical decision theoretical methods compared to using conventional, non-algorithmical selection procedure to arrive at the policy recommendation in form of optimal treatment (re-)allocation, Tables 5.2 and 5.3 enlist the key premia achieved by each method and its configuration. For comparability with preceding literature, I report these premia in form of *expected welfare gains* (see, e.g., welfare premia reported in Chapter 5 of Athey and Wager (2021). Here, the tables report the welfare gains (and the corresponding treated proportion) for the implementations on the SIAB and the BPJS-KIS health insurance datasets, respectively.

Examples of these welfare gains include the following. In Knaus (2022)'s use of Policy Learning for Swiss job training programs[38], the cumulated number of months in employment in the 31 months following the start of the program. As is common practice in the empirical literature implementing statistical decision theory up until now, Knaus (2022) reports the resulting decision trees as policy recommendations in (Figure 5, p. 620) – greatly similar to how I present the resulting decision trees for SIAB and BPJS-KIS settings in Figures 5.1–5.6 in Subchapter 5.4. The average increase of re-employment duration induced by each corresponding program is then reported in the columns 2–5 in Table 7 (Knaus 2022, p. 621) with the rows matching these results to its respective policy tree configuration (for example, using tree of depth 2 and 5 chosen variables, the job search program was shown to increase re-employment by 5.01 months; the result was lower with a configuration of depth 1 and 16 chosen variables at 4.73 months).

For the aforementioned Greater Avenus for Independence (GAIN) job training program in the United States starting in 1988, Athey and Wager (2021) reported

[38] Knaus (2022) studied over 100,000 cases of unemployment that were divided into non-participants as control group and a treatment group that in turn comprises four different programs: job search, vocational training, computer programs and language courses.

welfare gains in form of mean quarterly income averaged over a period of 9 years following the program assignment. Also analogously to my presentation in Subchapters 5.3–5.4, they reported the resulting decision trees from implementing their Policy Learning for the GAIN sample data each for the configurations with depth 1 and depth 2 (Figure 1, p. 154).

Following this standard practice in the literature strand, I report in Table 5.2 my estimates for average welfare gains in form of additional daily income (per day worked in the re-employment period; measured in Euro) in the SIAB German job training setting corresponding to each configuration of methods. First, following Kitagawa and Tetenov (2018) and Athey and Wager (2021), I report the average gain for the counterfactual "treat-all" case where everybody would have been assigned the aforementioned FbW program (hence proportion treated = 1) namely 39.62 € (daily income premium)[39].

The implementation of Empirical Welfare Maximization without blending out age and gender variables (i.e., as in the original method of Kitagawa and Tetenov, 2018) as benchmark yielded an average gain of 27.57 – 30% lower than the treat-all hypothetical case but with a trade-off of having to treat less than 5% of the sample population[40]. The average gain shown by implementing Policy Learning, on the other hand, jumps back nearing the treat-all average gain while choosing the treated proportion between of around 22–28.5% of the population – depending on the depth chosen.

Introducing the crucial feature of blending out politically/legally problematic variables (in SIAB case: age, gender) did not, interestingly, turn out to sacrifice

[39] As noted by Knaus (2022, p. 621), the current literature strand on statistical treatment assignment somewhat lacks a consensus on incorporating capacity constraints while maintaining the asymptotical guarantees proven in the literature as summarized in Athey and Wager (2021, Chapters 2–4).

[40] While I find the miniscule proportion of treated individuals chosen by the EWM method – especially when then compared to the proportions chosen by the Policy Learning method in its various depth configuration – to be utmost intriguing, decoding this choice would have necessitated further modification to the method along the line of interpretable machine learning discussed in Chapter 4 and thus currently methodologically out of the scope of this dissertation.

the resulting average gain – indeed, for each respective depth configuration the Policy Learning showed instead (slight) improvement, although the treatment proportion did rise for the depths 1 and 3 of the tree specifications. As hinted in Subchapter 5.3.2, the computationally significantly more complex depth-3 Policy Learning yielded virtually equal performance in terms of average gain as the simpler depth-1 – the latter achieving this at an even lower proportion of the treated. To point out once again, from the point of view of policy-relevance the depth-1 and depth-2 Policy Learning configurations are also preferrable than their depth-3 or deeper counterparts (cf., e.g., Athey and Wager 2021).

Similarly, I report in Table 5.3 my estimates for average welfare gains in form of efficiency gains (in the equivalent Indonesian Rupiahs referred to in the variables representing FKRTL and FKTP) from re-prioritizing the type of treatment (two-stages, with both FKTP and FKRTL present *or* one-stage, *directly* to the referrals to advance specialist treatment centers—the FKRTL). First, following Kitagawa and Tetenov (2018) and Athey and Wager (2021), I report the average gain for the counterfactual "treat-all" case where everybody would have been assigned a program (proportion treated = 1) – in this case 213,654 Indonesian Rupiahs (IDR).

The implementation of Empirical Welfare Maximization (Kitagawa and Tetenov, 2018) as benchmark method yielded an average gain of 109,524 IDR – over 50% lower than the treat-all hypothetical case and while treating virtually the full sample (at 0.99 treated proportion). The average gain shown by implementing Policy Learning with the default method choice of random forest at Step 1 is IDR 200,884 and IDR 189,893 for the depths 1 and 2 at the corresponding treatment proportion of 76.5% and 86%, respectively. A stronger performance was evident by the modified Policy Learning with neural network substituting random forest as prediction method in Step 1, with the resulting depth-1 decision tree delivering the strongest average gain of IDR 244,539 though at a higher treated proportion (92,5%).

Method	Tree depth	Average gain: additional daily income in €	Proportion treated
Treat-all	–	39.62	1
EWM	2	27.57	0.045
Policy learning	1	37.60	0.22
Policy learning	2	37.74	0.235
Policy learning	3	37.70	0.285
Policy learning (safe)	1	38.42	0.345
Policy learning (safe)	2	37.85	0.205
Policy learning (safe)	3	38.44	0.525

Table 5.2 *Average gain and treated proportion of each method configuration applied on SIAB dataset*

Method (*extensions*)	**Tree depth**	**Average gain: treatment cost reduction** in Indonesian Rupiahs (IDR)	**Proportion treated**
Treat-all	–	213,645	1
EWM	2	109,524	0.99
Policy learning (full Random Forest)	1	200,884	0.765
Policy learning (full Random Forest)	2	189,893	0.86
Policy learning (Step 1 with Neural Network)	1	244,539	0.925
Policy learning (Step 1 with Neural Network)	2	222,483	0.98

Table 5.3 Average gain and treated proportion of each method applied on BPJS-KIS health insurance dataset.

Chapter 5.6 Concluding Remarks

With regard to heterogeneity within the treated population and subsequently heterogeneous policy impact, Chapter 5 has fulfilled its threefold aim to (1) pioneer the use of state-of-the-art statistical and machine learning methods in calculating optimal policy allocation in the context of active labor market interventions in Germany and the health insurance, program in Indonesia; (2) extend the Empirical Welfare Maximization and Policy learning method from static, single-stage policy assignment into an empirical implementation with dynamic, multi-stage treatments; and (3) incorporate and assess further methodological extensions as portrayed in Chapter 4.4 for both empirical applications.

In sum, the empirical applications portrayed in this chapter further solidify the feasibility of improving program assignments using novel statistical decision methods such as EWM and Policy Learning. Future studies would particularly benefit from replicating the use of neural network (Deep Learning) in the first step as I did here in Subchapter 5.4, as well as complementing the variable importance and ALE plot with other interpretability methods. Another promising research direction would be to provide an in-depth comparison the results of using statistical treatment assignment in the fields of active job market policies (ALMP) and social insurance policies with the efficiency of conventional assignment in each respective sector. Finally, Knaus (2022) noted firstly that his particular dataset of Swiss active labor market program did not offer observed costs for each training program – a caveat shared by the version of SIAB dataset I used in Subchapter 5.3 –, but also more generally that the current currently available statistical decision theoretic methods have yet to develop how to incorporate capacity constraints while maintaining the statistical guarantees achieved by Policy Learning intact.

References

Abadie, A., & Cattaneo, M. D. (2018). Econometric Methods for Program Evaluation. *Annual Review of Economics*, *10*(1), 465–503. https://doi.org/10.1146/annurev-economics-080217-053402

Agersnap, O., & Zidar, O. (2021). The Tax Elasticity of Capital Gains and Revenue-Maximizing Rates. *American Economic Review: Insights*, *3*(4), 399–416. https://doi.org/10.1257/aeri.20200535

Angrist, J. D., & Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press. https://doi.org/10.1515/9781400829828

Antoni, Ganzer, & vom Berge. (2019). *Stichprobe der Integrierten Arbeitsmarktbiografien Regionalfile (SIAB-R) 1975-2017* (No. 04/2019; FDZ-Datenreport). IAB.

Ariawan, I., Sartono, B., & Wahyuningsih, W. (2020). *Data Sampel BPJS Kesehatan 2015-2018*.

Armstrong, T. B., & Shen, S. (2014). Inference on optimal treatment assignments. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2428925

Athey, S., & Imbens, G. W. (2017). The State of Applied Econometrics: Causality and Policy Evaluation. *Journal of Economic Perspectives*, *31*(2), 3–32. https://doi.org/10.1257/jep.31.2.3

Athey, S., & Imbens, G. W. (2019). Machine Learning Methods That Economists Should Know About. *Annual Review of Economics*, *11*(1), 685–725. https://doi.org/10.1146/annurev-economics-080217-053433

Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, *47*(2). https://doi.org/10.1214/18-AOS1709

Athey, S., & Wager, S. (2021). Policy learning with observational data. *Econometrica*, *89*(1), 40. https://doi.org/10.3982/ECTA15732

Atkinson, A. B., Piketty, T., & Saez, E. (2011). Top Incomes in the Long Run of History. *Journal of Economic Literature*, *49*(1), 3–71. https://doi.org/10.1257/jel.49.1.3

Auerbach, A. J., & Hines, J. R. (2002). Taxation and Economic Efficiency. In *Handbook of Public Economics* (Vol. 3, pp. 1347–1421). Elsevier. https://doi.org/10.1016/S1573-4420(02)80025-8

Bastani, S., & Waldenström, D. (2020). How should capital be taxed? *Journal of Economic Surveys*, *34*(4), 812–846. https://doi.org/10.1111/joes.12380

Bernhard, S. (2016). Berufliche Weiterbildung von Arbeitslosengeld-II-Empfängern. Langfristige Wirkungsanalysen. *Sozialer Fortschritt*, *65*(7), 153–161. https://doi.org/10.3790/sfo.65.7.153

Bhattacharya, D., & Dupas, P. (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, *167*(1), 168–196. https://doi.org/10.1016/j.jeconom.2011.11.007

Bock, J.-O., Heider, D., Matschinger, H., Brenner, H., Saum, K.-U., Haefeli, W. E., & König, H.-H. (2016). Willingness to pay for health insurance among the elderly population in Germany. *The European Journal of Health Economics*, *17*(2), 149–158. https://doi.org/10.1007/s10198-014-0663-8

Boersch-Supan, A., & Juerges, H. (2011). *Disability, Pension Reform and Early Retirement in Germany* (No. w17079; p. w17079). National Bureau of Economic Research. https://doi.org/10.3386/w17079

Boersch-Supan, A., & Schnabel, R. (1998). Social Security and Declining Labor-Force Participation in Germany. *American Economic Review*, *88*, 7.

Bohren, J. A., & Hauser, D. N. (2021). Learning With Heterogeneous Misspecified Models: Characterization and Robustness. *Econometrica*, *89*(6), 3025–3077. https://doi.org/10.3982/ECTA15318

Bundesinstitut für Berufsbildung. (2020). *Datenreport zum Berufsbildungsbericht 2020*. https://www.bibb.de/dienst/veroeffentlichungen/de/publication/show/16673

Busse, R., Blümel, M., Knieps, F., & Bärnighausen, T. (2017). Statutory health insurance in Germany: A health system shaped by 135 years of solidarity, self-governance, and competition. *The Lancet*, *390*(10097), 882–897. https://doi.org/10.1016/S0140-6736(17)31280-1

Caliendo, M., Hujer, R., & Thomsen, S. L. (2008). Identifying effect heterogeneity to improve the efficiency of job creation schemes in Germany. *Applied Economics*, *40*(9), 1101–1122. https://doi.org/10.1080/00036840500438897

Card, D., Kluve, J., & Weber, A. (2010). Active Labour Market Policy Evaluations: A Meta-Analysis. *The Economic Journal*, *120*(548), F452–F477. https://doi.org/10.1111/j.1468-0297.2010.02387.x

Card, D., Kluve, J., & Weber, A. (2018). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association*, *16*(3), 894–931. https://doi.org/10.1093/jeea/jvx028

Chendra, M. (2020). Pension reform in Indonesia: Issues and challenges. *Jurnal Institut BPJS Ketengakerjaan*, *5*(1), 75–107.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, *21*(1), C1–C68. https://doi.org/10.1111/ectj.12097

Chetty, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics*, *1*(1), 451–488. https://doi.org/10.1146/annurev.economics.050708.142910

Dehejia, R. H. (2005). Program evaluation as a decision problem. *Journal of Econometrics*, *125*(1–2), 141–173. https://doi.org/10.1016/j.jeconom.2004.04.006

Diamond, P., & Saez, E. (2011). The Case for a Progressive Tax: From Basic Research to Policy Recommendation. *Journal of Economic Perspectives*, *25*(4), 165–190. https://doi.org/10.1257/jep.25.4.165

Doerrenberg, P., Peichl, A., & Siegloch, S. (2017). The elasticity of taxable income in the presence of deduction possibilities. *Journal of Public Economics*, *151*, 41–55. https://doi.org/10.1016/j.jpubeco.2015.10.001

Einav, L., & Finkelstein, A. (2023). Empirical analyses of selection and welfare in insurance markets: A self-indulgent survey. *The Geneva Risk and Insurance Review*, *48*(2), 167–191. https://doi.org/10.1057/s10713-023-00084-3

Einav, L., Finkelstein, A., & Cullen, M. R. (2010). Estimating Welfare in Insurance Markets Using Variation in Prices*. *Quarterly Journal of Economics*, *125*(3), 877–921. https://doi.org/10.1162/qjec.2010.125.3.877

Finkelstein, A., & Hendren, N. (2020). Welfare Analysis Meets Causal Inference. *Journal of Economic Perspectives*, *34*(4), 146–167. https://doi.org/10.1257/jep.34.4.146

Finkelstein, A., Hendren, N., & Shepard, M. (2019). Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts. *American Economic Review*, *109*(4), 1530–1567. https://doi.org/10.1257/aer.20171455

Fitzenberger, B., & Speckesser, S. (2007). Employment effects of the provision of specific professional skills and techniques in Germany. *Empirical Economics*, *32*(2–3), 529–573. https://doi.org/10.1007/s00181-006-0088-z

García, J. L., & Heckman, J. (2022a). *Three Criteria for Evaluating Social Programs* (No. w30507; p. w30507). National Bureau of Economic Research. https://doi.org/10.3386/w30507

García, J. L., & Heckman, J. J. (2022b). *On Criteria for Evaluating Social Programs* (No. 30005; NBER Working Paper, p. 18).

García-Miralles, E. (2020). The Crucial Role of Social Welfare Criteria and Individual Heterogeneity for Optimal Inheritance Taxation. *The B.E. Journal of Economic Analysis & Policy*, *20*(2), 20190274. https://doi.org/10.1515/bejeap-2019-0274

Gesundheitsberichterstattung des Bundes. (2021). *Gesundheitsausgaben in Deutschland*. Statistisches Bundesamt. https://www.gbe-bund.de/gbe/pkg_olap_tables.prc_set_hierlevel?p_uid=gast&p_aid=76630471&p_sprache=D&p_help=2&p_indnr=863&p_ansnr=51503225&p_version=2&p_dim=D.733&p_dw=4438&p_direction=drill

Glogowsky, U. (2021). Behavioral responses to inheritance and gift taxation: Evidence from Germany. *Journal of Public Economics*, *193*, 104309. https://doi.org/10.1016/j.jpubeco.2020.104309

Görlitz, K., & Tamm, M. (2015). The Pecuniary and Non-Pecuniary Returns to Voucher-Financed Training. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2618444

Gottfried, & Witczak. (2009). The responses of taxable income induced by tax cuts: Empirical evidence from the German taxpayer panel. *IAW-Diskussionspapiere*, *57*. http://hdl.handle.net/10419/39211

Goupille-Lebret, J., & Infante, J. (2018). Behavioral responses to inheritance tax: Evidence from notches in France. *Journal of Public Economics*, *168*, 21–34. https://doi.org/10.1016/j.jpubeco.2018.09.016

Hajek, A., Enzenbach, C., Stengler, K., Glaesmer, H., Hinz, A., Röhr, S., Stein, J., Riedel-Heller, S. G., & König, H.-H. (2020). Determinants of Willingness to Pay for Health Insurance in Germany—Results of the Population-Based Health Study of the Leipzig Research Centre for Civilization Diseases (LIFE-Adult-Study). *Frontiers in Public Health*, *8*, 456. https://doi.org/10.3389/fpubh.2020.00456

Harrer, T., & Stockinger, B. (2019). Ein-Euro-Jobs nach der Instrumentenreform 2012: Zielgruppe besser erreicht – erste Ergebnisse zur Wirkung. *IAB-Kurzbericht*, *22*, 8.

Hendren, N. (2016). The Policy Elasticity. *Tax Policy and the Economy*, *30*(1), 51–89. https://doi.org/10.1086/685593

Hendren, N. (2021). Measuring Ex Ante Welfare in Insurance Markets. *The Review of Economic Studies*, *88*(3), 1193–1223. https://doi.org/10.1093/restud/rdaa015

Hendren, N., Landais, C., & Spinnewijn, J. (2021). Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design. *Annual Review of Economics*, *13*(1), 457–486. https://doi.org/10.1146/annurev-economics-090820-111716

Hendren, N., & Sprung-Keyser, B. (2020). A Unified Welfare Analysis of Government Policies*. *The Quarterly Journal of Economics*, *135*(3), 1209–1318. https://doi.org/10.1093/qje/qjaa006

Hendren, N., & Sprung-Keyser, B. (2022). *The Case for Using the MVPF in Empirical Welfare Analysis* (No. 30029; NBER Working Paper, p. 11). https://scholar.harvard.edu/files/hendren/files/mvpf_case_vfinal.pdf

Herberholz, L., & Wigger, B. U. (2021). Efficiency of European universities: A comparison of peers. *Research Policy*, *50*(9), 104314. https://doi.org/10.1016/j.respol.2021.104314

Hirano, K., & Porter, J. R. (2009). Asymptotics for statistical treatment rules. *Econometrica*, *77*(5), 1683–1701. https://doi.org/10.3982/ECTA6630

Hirano, K., & Porter, J. R. (2020). Asymptotic analysis of statistical decision rules in econometrics. In *Handbook of Econometrics* (p. S1573441220300040). Elsevier. https://doi.org/10.1016/bs.hoe.2020.09.001

Hirshleifer, J. (1971). The Private and Social Value of Information and the Reward to Inventive Activity. *American Economic Review*, *61*(4), 561–574.

Huber, M., Lechner, M., Wunsch, C., & Walter, T. (2011). Do German Welfare-to-Work Programmes Reduce Welfare Dependency and Increase Employment? *German Economic Review*, *12*(2), 182–204. https://doi.org/10.1111/j.1468-0475.2010.00515.x

Jakobsen, K., Jakobsen, K., Kleven, H., & Zucman, G. (2020). Wealth Taxation and Wealth Accumulation: Theory and Evidence From Denmark*. *The Quarterly Journal of Economics*, *135*(1), 329–388. https://doi.org/10.1093/qje/qjz032

Kasy, M. (2018). Optimal taxation and insurance using machine learning—Sufficient statistics and beyond. *Journal of Public Economics*, *167*, 205–219. https://doi.org/10.1016/j.jpubeco.2018.09.002

Kemper, L., Vorhoff, G., & Wigger, B. U. (2020). Predicting student dropout: A machine learning approach. *European Journal of Higher Education*, *10*(1), 28–47. https://doi.org/10.1080/21568235.2020.1718520

Kitagawa, T., & Tetenov, A. (2018). Who should be treated? Empirical Welfare Maximization methods for treatment choice. *Econometrica*, *86*(2), 591–616. https://doi.org/10.3982/ECTA13288

Kleven, H. J. (2021). Sufficient Statistics Revisited. *Annual Review of Economics*, *13*(1), 515–538. https://doi.org/10.1146/annurev-economics-060220-023547

Kleven, H. J., & Kreiner, C. T. (2006). The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics*, *90*(10–11), 1955–1973. https://doi.org/10.1016/j.jpubeco.2006.03.006

Kleven, H. J., & Schultz, E. A. (2014). Estimating Taxable Income Responses Using Danish Tax Reforms. *American Economic Journal: Economic Policy*, *6*(4), 271–301. https://doi.org/10.1257/pol.6.4.271

Klös, H.-P. (2021). *Berufliche Weiterbildung in Deutschland: Status Quo und Weiterentwicklung*. 42.

Kluve, J. (2010). The effectiveness of European active labor market programs. *Labour Economics*, *17*(6), 904–918. https://doi.org/10.1016/j.labeco.2010.02.004

Knaus, M. C. (2022). Double machine learning-based programme evaluation under unconfoundedness. *The Econometrics Journal*, *25*(3), 602–627. https://doi.org/10.1093/ectj/utac015

Kopczuk, W. (2013). Taxation of Intergenerational Transfers and Wealth. In *Handbook of Public Economics* (Vol. 5, pp. 329–390). Elsevier. https://doi.org/10.1016/B978-0-444-53759-1.00006-6

Krolage, C. (2020). *Essays in Empirical Public Economics*. LMU.

Landais, C., Nekoei, A., Nilsson, P., Seim, D., & Spinnewijn, J. (2021). Risk-Based Selection in Unemployment Insurance: Evidence and Implications.

American Economic Review, *111*(4), 1315–1355. https://doi.org/10.1257/aer.20180820

Le Barbanchon, T., Schmieder, J., & Weber, A. (2024). Job search, unemployment insurance, and active labor market policies. In *Handbook of Labor Economics* (Vol. 5, pp. 435–580). Elsevier. https://doi.org/10.1016/bs.heslab.2024.11.006

Lechner, M. (2023). Causal Machine Learning and its use for public policy. *Swiss Journal of Economics and Statistics*, *159*(1), 8. https://doi.org/10.1186/s41937-023-00113-y

Lechner, M., Miquel, R., & Wunsch, Conny. (2011). Longrun effects of public sector sponsored training in West Germany. *Journal of the European Economic Association*, *9*(4), 742–784.

Lefebvre, M., Lehmann, E., & Sicsic, M. (2025). Estimating the Laffer tax rate on capital income: Cross-base responses matter! *The Scandinavian Journal of Economics*, *127*(2), 460–489. https://doi.org/10.1111/sjoe.12578

Löffler, M., & Siegloch, S. (2021). *Welfare Effects of Property Taxation*.

Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, *72*(4), 1221–1246. https://doi.org/10.1111/j.1468-0262.2004.00530.x

Manski, C. F. (2021). Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald. *Econometrica*, *89*(6), 2827–2853. https://doi.org/10.3982/ECTA17985

Mayshar, J. (1990). On measures of excess burden and their application. *Journal of Public Economics*, *43*(3), 263–289. https://doi.org/10.1016/0047-2727(90)90001-X

Mbakop, E., & Tabord-Meehan, M. (2021). Model Selection for Treatment Choice: Penalized Welfare Maximization. *Econometrica*, *89*(2), 825–848. https://doi.org/10.3982/ECTA16437

McCall, B., Smith, J., & Wunsch, C. (2016). Government-Sponsored Vocational Education for Adults. In *Handbook of the Economics of Education* (Vol. 5, pp. 479–652). Elsevier. https://doi.org/10.1016/B978-0-444-63459-7.00009-9

Molnar, C. (2022). *Model-agnostic interpretable machine learning* [Ludwig-Maximilians-Universität München; Application/pdf]. https://doi.org/10.5282/EDOC.30374

Pigou, A. C. (1920). *The Economics of Welfare*.

Piketty, T., & Saez, E. (2013). A Theory of Optimal Inheritance Taxation. *Econometrica*, *81*(5), 1851–1886. https://doi.org/10.3982/ECTA10712

Piketty, T., Saez, E., & Zucman, G. (2023). *Rethinking Capital and Wealth Taxation* (No. halshs-04104410). Hal SHS. https://shs.hal.science/halshs-04104410

Saez, E., Slemrod, J., & Giertz, S. H. (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, *50*(1), 3–50. https://doi.org/10.1257/jel.50.1.3

Saez, E., & Stantcheva, S. (2018). A simpler theory of optimal capital taxation. *Journal of Public Economics*, *162*, 120–142. https://doi.org/10.1016/j.jpubeco.2017.10.004

Schmidt, T.-P., & Müller, H. (2012). Die Elastizität des zu versteuernden Einkommens in Deutschland. *Arqus Discussion Paper*, *142*, 36. http://dx.doi.org/10.2139/ssrn.2088106

Schmieder, J. F., & von Wachter, T. (2017). A Context-Robust Measure of the Disincentive Cost of Unemployment Insurance. *American Economic Review*, *107*(5), 343–348. https://doi.org/10.1257/aer.p20171026

Schwerdt, G., Messer, D., Woessmann, L., & Wolter, S. C. (2012). The impact of an adult education voucher program: Evidence from a randomized field experiment. *Journal of Public Economics*, *96*(7–8), 569–583. https://doi.org/10.1016/j.jpubeco.2012.03.001

Setio, E. S. (2021). Comparative Welfare Analysis of German Public Policies. *Journal of Business and Political Economy: Biannual Review of The Indonesian Economy*, *3*(2), 125–147.

Siegloch, S., Wehrhöfer, N., & Etzel, T. (2021). Direct, Spillover and Welfare Effects of Regional Firm Subsidies. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3839263

Slemrod, J., & Yitzhaki, S. (1996). The Costs of Taxation and the Marginal Efficiency Cost of Funds. *Staff Papers - International Monetary Fund*, *43*(1), 172. https://doi.org/10.2307/3867356

Slemrod, J., & Yitzhaki, S. (2001). Integrating Expenditure and Tax Decisions: The Marginal Cost of Funds and the Marginal Benefit of Projects. *National Tax Journal*, *54*(2), 189–201. https://doi.org/10.17310/ntj.2001.2.01

Slesnick, D. T. (1998). Empirical Approaches to the Measurement of Welfare. *Journal of Economic Literature*, *36*(4), 59. https://www.jstor.org/stable/2565048

Spinnewijn, J. (2020). The Trade-Off between Insurance and Incentives in Differentiated Unemployment Policies. *Fiscal Studies*, *41*(1), 101–127. https://doi.org/10.1111/1475-5890.12199

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics*, *151*(1), 70–81. https://doi.org/10.1016/j.jeconom.2009.02.013

Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics*, *166*(1), 138–156. https://doi.org/10.1016/j.jeconom.2011.06.012

Vooren, M., Haelermans, C., Groot, W., & Maassen van den Brink, H. (2019). The effectiveness of active labor market policies: A meta-analysis: the effectiveness of active labor market policies: a meta-analysis. *Journal of Economic Surveys*, *33*(1), 125–149. https://doi.org/10.1111/joes.12269

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

Wald, A. (1950). *Statistical decision functions.* (pp. ix, 179). Wiley.

Werdt, C. (2015). *The elasticity of taxable income for Germany and its sensitivity to the appropriate model* (No. 2015/5; Diskussionsbeiträge). http://hdl.handle.net/10419/106825

Ye, H. (2018). *Essays on the economics of unemployment and retirement* [Thesis (PhD)]. Boston University.