Carbon Pricing and Redistribution*

Lea Fricke
University of St. Gallen

Clemens Fuest
ifo Institute, University of Munich

Dominik Sachs
University of St. Gallen

August 13, 2026

Abstract

We show that carbon pricing fundamentally alters the equity-efficiency trade-off, shifting optimal policy toward greater redistribution. We study optimal redistribution and carbon pricing in an optimal income taxation model with non-homothetic demand, where households differ in carbon footprints through both spending levels and the composition of their consumption baskets. Carbon pricing separates the effective marginal tax rate from the welfare-relevant labor wedge: labor supply reductions reduce not only tax revenue, but also emissions. This lowers the welfare cost of redistribution. The optimal policy therefore increases redistribution, but not as much as a lump-sum carbon dividend would. Instead, the optimal allocation is implemented by a rebate schedule that rises with income. In a quantitative application to Germany, the optimal rebate at the 90th income percentile is nearly three times that at the 10th percentile, whereas carbon tax payments are about eight times as large. The resulting carbon policy is therefore strongly progressive.

JEL codes: H21, H23, Q58

Keywords: Carbon Tax, Optimal Taxation

*lea.fricke@student.unisg.ch, fuest@ifo.de and dominik.sachs@unisg.ch. Parts of this research were conducted while Lea Fricke was a visiting researcher at the ifo Institute. We thank Paul Ruwe and Anna Zeitz for excellent research assistance. We thank Marcelo Arbex, Johannes Becker, Sabrina Eisenbarth, Xavier Jaravel, Albert Jan Hummel, Louis Kaplow, Winfried Koeniger, Dirk Krueger, Andrea Lanteri, Etienne Lehmann, Arash Nekoei, Casey Rothschild, Florian Scheuer, Lennart Schlattmann, Karl Schulz and Daniel Weishaar for helpful comments. We are grateful to the LMU-ifo Economics Business Data Center (EBDC), and especially to Michael Rindler, for facilitating access to the data. We further thank seminar and conference participants at IIPF 2024, Swiss Macro Workshop 2025, CY Cergy Paris, 2025 Meeting of the Public Finance Committee of the German Economic Association, ZEW Public Finance Conference 2025, Barcelona Workshop on the Economics of Taxation 2025, CESifo Public Economics Week 2025, KU Leuven, CESifo Area Conference on Macro with Micro Data 2026, Bristol Macro Workshop 2026 and NBER SI Macro Public Finance 2026 for comments.

1 Introduction

Carbon pricing is a central instrument for reducing greenhouse gas emissions (Barrage and Nordhaus 2024), but it is politically contested because it creates a distributional conflict. The Yellow Vests protests in France illustrate the broader challenge: higher prices for fuel, heating, and other carbon-intensive necessities reduce real purchasing power, and because lower-income households devote larger budget shares to such goods, the burden can fall relatively more heavily on them (Grainger and Kolstad 2010, Bigio, Känzig, Sánchez, and Walsh 2026). Compensation is therefore widely viewed as central to socially and politically sustainable carbon pricing (Blanchard, Gollier, and Tirole 2023), and survey evidence suggests that support depends not only on the carbon price itself, but also on who bears the burden and how revenues are rebated (Dechezleprêtre, Fabre, Kruse, Planterose, Sanchez Chico, and Stantcheva 2025).

This paper asks how fiscal policy should respond to the distributional conflict created by carbon pricing. Our central result is that carbon pricing does not merely create a compensation problem within the standard equity-efficiency trade-off; it relaxes the trade-off itself. When output generates carbon emissions, reductions in labor supply lower output and therefore reduce emissions. Redistribution therefore becomes less costly in welfare terms, tilting optimal policy toward greater equity and more redistribution, holding the welfare criterion fixed.[1]

We formalize and quantify this mechanism in an optimal income taxation model with carbon emissions and non-homothetic demand. The strength of the mechanism is governed by carbon inequality: the distribution of carbon footprints and carbon tax burdens across the income distribution. This income gradient matters for incidence and incentives. For incidence, the income gradient in the carbon tax burdens determines how the costs of carbon pricing are distributed across the income distribution. For incentives, the local derivative of emissions with respect to net expenditure – the *marginal propensity to pollute* – determines how labor supply responses translate into emissions changes. The larger this marginal propensity to pollute, the more strongly redistribution reduces emissions and the larger the optimal increase in redistribution.

A useful benchmark is a lump-sum carbon dividend, which pools carbon tax revenue and pays the same rebate to every household, leaving the income gradient in carbon tax liabilities fully reflected in net burdens. The optimal policy generally does not go this far. Instead, it offsets part of the income gradient in carbon tax liabilities by making rebates rise with income. Nevertheless, low-income households are overcompensated relative to their carbon tax burden and high-income households are undercompensated. The overall policy is more redistributive than the status quo but less redistributive than a carbon dividend.

[1] This mechanism is akin to the basic intuition behind the degrowth debate: as long as additional economic activity is associated with carbon emissions, reductions in output carry an environmental benefit.

We quantify the model for Germany using household expenditure microdata and input-output based carbon emission intensities. Carbon inequality is large: households at the 90th income percentile pay roughly eight times as much in carbon taxes as households at the 10th percentile, while the optimal rebate is only about three times as large. Despite the upward-sloping rebate schedule, the resulting carbon policy is therefore progressive.

Theory. We consider an economy in which households differ by skill, make labor supply and consumption decisions, and have non-homothetic preferences that generate heterogeneous consumption baskets across the income distribution. Carbon emissions arise from the consumption of goods with different emission intensities. We compare a status-quo economy in which the government does not account for carbon emissions and sets a nonlinear income tax to balance equity and efficiency with an otherwise identical economy in which the government additionally faces a binding carbon cap. To meet the emissions target, the government imposes a uniform carbon tax and adjusts the income tax schedule accordingly. The resulting change in the income tax schedule can be interpreted as an income-dependent carbon rebate. This setup isolates the effect of carbon constraints on optimal redistribution in a clean ceteris paribus comparison. In particular, we deliberately consider an environment in which the Atkinson–Stiglitz assumptions hold and the optimal carbon tax therefore coincides with the Pigouvian benchmark. Thus, we shut down the possibility that distributional considerations move the carbon tax itself away from the Pigouvian level and isolate instead how optimal redistribution changes once emissions are efficiently priced.[2]

Our central analytical result is that carbon pricing separates two objects that coincide in the standard benchmark. Absent a binding carbon constraint, the effective marginal tax rate plays two roles: it is the local measure of marginal redistribution and it coincides with the welfare-relevant labor wedge, the local measure of the efficiency cost of labor supply distortions. Under carbon pricing, by contrast, these two objects diverge. The effective marginal tax rate continues to measure marginal redistribution, but the welfare-relevant labor wedge is smaller. The reason is that labor supply reductions not only reduce tax revenue; they also reduce emissions and relax the carbon constraint. We call the resulting gap between the effective marginal tax rate and the welfare-relevant labor wedge the *carbon redistribution wedge*. It is the local measure of the redistributive scope created by carbon pricing. The wedge is larger when the carbon tax is higher and when the local carbon-income gradient, as captured by the marginal propensity to pollute, is steeper. Estimating this gradient therefore quantifies the additional redistributive scope created by carbon pricing.

[2]For conditions under which the optimal environmental tax follows the Pigouvian rule in the presence of redistributive income taxation, see Kaplow (2012) and Jacobs and De Mooij (2015). A related literature studies when distributional considerations or other forms of heterogeneity lead to departures from this benchmark; see Jacobs and van der Ploeg (2019), Hänsel, Franks, Kalkuhl, and Edenhofer (2022), Craig, Lloyd, and Moore (2025), Ahlvik, Liski, and Mäkimattila (2025), and Bierbrauer (2025). We discuss this literature in detail below.

This decomposition clarifies the role of alternative rebate schemes. A distributionally neutral rebate offsets each household's carbon tax liability and leaves effective marginal tax rates unchanged. It therefore leaves the new redistributive scope unused. A lump-sum carbon dividend is the opposite benchmark. It does not offset the income gradient in carbon tax liabilities, so the marginal carbon tax burden is fully reflected in higher effective marginal tax rates. In this sense, the carbon dividend exhausts the redistributive scope created by the carbon constraint. Our comparative statics suggest that the optimum lies between these two benchmarks, with its precise location depending on how quickly marginal distributional gains decline as redistribution increases. If these gains decline strongly, the redistributive adjustment is more limited. If they decline only weakly, policy moves closer to the carbon dividend. In the polar case in which marginal distributional gains are constant, the carbon dividend is optimal. The quantitative analysis confirms this intermediate pattern: the optimal policy increases redistribution relative to the status quo, but by less than a carbon dividend would. Equivalently, the optimal rebate schedule rises with income, but less steeply than carbon tax payments.

Measurement and quantitative framework. Quantifying the mechanism requires two empirical objects: the level and the local slope of the carbon Engel curve. The former governs carbon tax liabilities across the income distribution; the latter, the marginal propensity to pollute, governs the carbon redistribution wedge. We construct both objects for Germany in 2018. We combine detailed expenditure microdata for 107 COICOP goods and services with consumption-based emission intensities that account for direct household emissions and indirect emissions throughout global production chains. This yields percentile-level carbon footprints and marginal propensities to pollute along the gross labor income distribution.

The carbon Engel curves alone are not sufficient for counterfactual carbon pricing, because households also respond to the induced changes in relative prices. We therefore fit a two-tier non-homothetic CES demand system. At the upper tier, households allocate expenditure across nine broad consumption categories. Within each category, a second non-homothetic CES nest allocates expenditure between brown and green varieties. The model is disciplined by expenditure patterns along the income distribution and by product-level emission intensities. It captures changes in consumption composition at both margins and allows carbon pricing to induce substitution both across categories and between cleaner and dirtier varieties within categories.

We combine this demand system with a DICE-style abatement technology, administrative income-tax data, and the estimated German tax-transfer system. The resulting framework links granular consumption-based carbon accounting to a structural model of demand, abatement, and labor supply, allowing us to quantify both the incidence of carbon pricing and the strength of the carbon redistribution wedge. We hold fixed the welfare criterion that rationalizes the

2018 tax-transfer system across all counterfactuals, thereby isolating how a binding carbon constraint changes optimal redistribution.

Quantitative results. We first evaluate a carbon cap that reduces aggregate emissions by 25%. The optimal rebate schedule is increasing in income: households at the 90th income percentile receive 3.31 times the rebate of households at the 10th percentile. Yet the overall carbon policy is strongly progressive, because carbon tax payments rise much more steeply with income: households at the 90th percentile pay 7.87 times as much in carbon taxes as households at the 10th percentile. We provide a broad set of robustness checks across alternative parameter values, and the main quantitative results are remarkably stable across specifications.

We then trace the optimal policy along the decarbonization path by varying the emission reduction target. The relative incidence patterns remain remarkably stable: the 90-to-10 ratios of optimal rebates and carbon tax payments stay close to their benchmark values. What changes is the quantitative importance of the policy. As the carbon cap tightens, carbon tax revenue initially rises, so the induced increase in redistribution becomes larger. This effect peaks around a 65% emissions reduction, when the increase in the gap between the effective average tax rates at the 10th and 90th percentiles reaches about 4 percentage points and carbon tax revenue is equivalent to that generated by an almost 5 percentage point increase in the statutory VAT rate. Beyond that point, the shrinking carbon tax base dominates and revenues begin to decline. Thus, compensation and rebate design become increasingly important over a substantial part of the decarbonization path, not only for small carbon pricing reforms.

Related literature. The interaction between income and carbon taxation was first analyzed in the so-called double-dividend literature. In representative agents models, it has been shown that the optimal carbon tax should be below the Pigouvian tax when distortionary income taxes are present, since carbon taxation erodes the income tax base (see, e.g., Bovenberg and De Mooij, 1994; Bovenberg and van der Ploeg, 1994; Goulder, 1995). A recent modern quantitative treatment of this channel can be found in Barrage (2020) who finds that the optimal carbon tax should be 8-24% lower.

Kaplow (2012) and Jacobs and De Mooij (2015) integrate externalities into a framework with heterogeneous agents and redistributive income taxation.[3] Both show that, under weak separability, the optimal carbon tax coincides with the Pigouvian tax, even in the presence of distortionary income taxation. Recent complementary work by Craig, Lloyd, and Moore (2025) asks when distributional concerns justify moving environmental taxes away from the Pigouvian benchmark. They show that the Pigouvian benchmark can fail when differences in

[3]Focusing on a linear income tax, Jacobs and van der Ploeg (2019) show that if carbon Engel curves are linear, the optimal pollution tax follows a first-best rule. With non-linear carbon Engel curves, however, carbon taxes acquire a redistributive role that may push them above or below the Pigouvian level. Hänsel, Franks, Kalkuhl, and Edenhofer (2022) extend the analysis by adding horizontal inequality in carbon emissions and by considering subsidies for clean energy as an additional policy instrument.

carbon intensity across the income distribution are driven by preference heterogeneity rather than by pure income effects. Our paper complements this literature. We deliberately focus on the weak-separability environment in which the optimal carbon tax continues to equal the Pigouvian benchmark. We show that even in that setting, carbon pricing changes the equity-efficiency trade-off itself and thereby the optimal degree of redistribution as well as the schedule of income-dependent rebates.[4]

Our mechanism is closely related to the tax-interaction literature, especially the work of Goulder. This literature shows how environmental taxes interact with pre-existing labor taxes and how recycling revenues through lower marginal tax rates can mitigate these effects (Goulder 1995, Goulder 1998, Goulder 2013, Goulder, Parry, Williams III, and Burtraw 1999). We recast this mechanism in a Mirrleesian setting with heterogeneous households and nonlinear income taxation. At the optimal carbon price, the carbon tax revenue component of the fiscal externality is exactly offset by the welfare cost of the associated emissions, creating a pointwise gap between the effective marginal tax rate and the welfare-relevant labor wedge.

More broadly, our mechanism relates to work on optimal income taxation when higher earnings are associated with negative externalities through relative-income or status concerns (Boskin and Sheshinski 1978, Ireland 2001, Frank 1985). Rothschild and Scheuer (2016) derive an externality-corrective component in the nonlinear income tax when income can arise from rent-seeking; their broader multi-activity framework in Rothschild and Scheuer (2014) allows for general patterns of positive and negative externalities.

The crucial difference concerns the scope for directly targeting the externality. In our setting, emissions can be directly targeted through a carbon tax, while income taxation alone would not suffice because it cannot induce substitution toward cleaner goods conditional on expenditure. The relevant role of the income tax adjustment is therefore not to correct the externality itself, but to determine how much of the carbon tax induced increase in marginal tax payments is offset. A distributionally neutral rebate would offset it fully; we show that the optimal rebate does not, so effective marginal tax rates can rise while statutory marginal income-tax rates remain unchanged or fall.

This distinction connects our results directly to the literature on additivity and targeting. Sandmo (1975) established the classic additivity result: with an atmospheric externality, the externality alters the optimal tax on the externality generating good while leaving the formulas for other tax instruments unchanged. Cremer, Gahvari, and Ladoux (1998) show that whether this logic extends to income taxation depends on the available commodity tax instruments, while Kopczuk (2003) formulates a more general principle of targeting. Micheletto (2008) shows

[4] Ahlvik, Liski, and Mäkimattila (2025) study a Mirrleesian environment with multidimensional heterogeneity in earning ability and mitigation costs. They show that the optimal externality tax varies with income and apply their framework to vehicle use and electricity consumption. Their analysis focuses on income as a signal of heterogeneous mitigation costs and horizontal-equity concerns, whereas our mechanism operates through the effect of additional expenditure on emissions. Bierbrauer (2025) studies, more generally, conditions under which a uniform carbon price is optimal in the presence of heterogeneity and alternative production technologies.

that additivity can fail for non-atmospheric consumption externalities, including positional interpretations, and derives conditions under which it is restored. Consistent with additivity, the optimal statutory marginal income tax condition retains the standard form. Carbon pricing nevertheless changes effective redistribution: the marginal carbon tax burden raises the effective marginal tax rate without raising the welfare-relevant labor wedge, and the optimal rebate schedule does not generally offset this increase in full.

Our paper also relates to recent quantitative work on carbon pricing and heterogeneous households. Goulder, Hafstead, Kim, and Long (2019) quantify the distributional incidence of carbon taxes across US households and show that the net burden depends crucially on how revenues are recycled. van der Ploeg, Rezai, and Tovar Reaños (2022) study how alternative recycling schemes shape equity and political support for green tax reform using German household survey data. Douenne, Hummel, and Pedroni (2026) analyze jointly optimal carbon and income taxes in a dynamic climate-economy model with heterogeneous households and Douenne, Dyrda, Hummel, and Pedroni (2025) extend the analysis to incomplete markets and uninsurable income risk. By contrast, our paper makes explicit and quantifies a mechanism through which carbon pricing affects redistribution: it opens a gap between the effective marginal tax rate and the welfare-relevant labor wedge.

On the empirical side, our paper contributes to the literature that constructs consumption-based household carbon footprints by combining household expenditure micro data with environmentally extended input-output accounts. These measures have in turn been used to estimate environmental Engel curves, tracing how household emissions vary with income or expenditure (Wier, Lenzen, Munksgaard, and Smed 2001, Weber and Matthews 2008, Levinson and O'Brien 2019, Sager 2019, Hardadi, Buchholz, and Pauliuk 2021 among others). We provide a detailed and transparent implementation for Germany, making explicit the construction of direct and indirect product-level emission intensities and their mapping to household expenditure categories.

Structure. In Section 2, we explain the main mechanism for our result in a simple model. We then introduce our model in Section 3 and characterize the optimal tax-transfer system in the absence of carbon policies in Section 4. In Section 5 we then turn to the optimal design of carbon rebate policies. In Section 6, we construct the empirical inputs and calibrate the quantitative model to the German economy in 2018. Section 7 presents the quantitative results. Section 8 concludes.

2 The Main Mechanism

This section provides a reduced-form explanation of the mechanism underlying our main result. Consider a representative household of productivity θ who chooses labor supply l. Gross labor income is then given by $y = l\theta$ and expenditure equals $e(y) = y - \mathcal{T}(y)$, where $\mathcal{T}(y)$ is the tax-transfer system. Consumption expenditure e generates carbon emissions $f(e)$, which are taxed at rate τ_{co2} per unit of carbon. Total tax revenue raised from this household is given by:

$$\mathcal{R}(y) = \mathcal{T}(y) + \tau_{\text{co2}} f(e(y)).$$

Effective marginal tax rate. Define the effective marginal tax rate (EMTR) as the increase in fiscal revenue induced by a one-unit increase in income:

$$\tau_{\text{eff}}(y) \equiv \frac{d\mathcal{R}(y)}{dy} = \mathcal{T}'(y) + \tau_{\text{co2}}(1 - \mathcal{T}'(y))MPP(y). \tag{1}$$

where $\mathcal{T}'(y)$ is the marginal income tax rate and $MPP(y)$ is the marginal propensity to pollute out of expenditure of an individual with income y, formally defined by $MPP(y) \equiv f'(e(y))$. It captures the increase in the carbon footprint that results from one more unit of spending. Hence, the EMTR exceeds the marginal income tax rate whenever higher expenditure raises emissions, because an additional unit of gross income generates not only income tax revenue but also carbon tax revenue. In general, a higher $\tau_{\text{eff}}(y)$ means that the total tax burden rises more steeply with gross income, implying that the tax system is more redistributive at the margin.

In the absence of externalities, the EMTR is the relevant metric for the welfare effects of changes in labor supply in response to policies. By the envelope theorem, a change in labor supply has no first-order effects on individual utility. To first order, the welfare effect of a change in labor supply operates through the induced change in government revenue (fiscal externalities) equal to $dy \times \tau_{\text{eff}}(y)$ (see e.g. Hendren (2016)).

Welfare wedge. In the presence of externalities, the EMTR no longer summarizes the welfare effects of labor supply responses because an income change also affects carbon emissions. To make this explicit, consider a government that maximizes a standard welfare function $\mathcal{W}$ subject to (i) a revenue requirement and (ii) a carbon emission cap $F \leq \bar{F}$, where F denotes aggregate emissions in the economy and $\bar{F}$ is an exogenous limit on emissions.[5] Let $\mathcal{L}$ denote the Lagrangian of the government optimization problem, and let Λ and μ be the multipliers on the revenue requirement and the emissions cap, respectively.

Consider the welfare effect of a marginal increase in gross labor income y induced by some policy change:

$$\frac{\partial \mathcal{L}}{\partial y} dy = \Lambda \left(\underbrace{\tau_{\text{eff}}(y)}_{\text{fiscal externality}} - \underbrace{\frac{\mu}{\Lambda}(1 - \mathcal{T}'(y))MPP(y)}_{\text{carbon cap externality}} \right) dy \equiv \Lambda \underbrace{\mathrm{T}(y)}_{\text{welfare wedge}} dy. \tag{2}$$

[5]Equivalently, one could incorporate the atomistic externality directly into the welfare function. The argument would be unchanged, see Section 3.3.

where we use $df/dy = (1-\mathcal{T}'(y))MPP(y)$. The first term in (2) captures the fiscal gain from higher labor income, while the second term captures the welfare loss from tightening the emissions constraint. The welfare wedge $\mathrm{T}(y)$ is, by construction, the welfare-relevant wedge for labor supply changes:

$$\mathrm{T}(y) = \mathcal{T}'(y) + \left(\tau_{\text{co2}} - \frac{\mu}{\Lambda}\right)(1-\mathcal{T}'(y))\, MPP(y).$$

Two implications follow immediately. First, if the carbon cap is slack ($\mu = 0$), then we obtain the usual result that this wedge equals the EMTR: $\mathrm{T}(y) = \tau_{\text{eff}}(y)$. In this case, the welfare effect of an income change is fully summarized by the EMTR. Second, if the cap binds ($\mu > 0$) and emissions rise with income ($MPP(y) > 0$), then $\mathrm{T}(y) < \tau_{\text{eff}}(y)$. In this case, a labor supply reduction has the positive side effect that it relaxes the carbon cap constraint.

As we show formally in Lemma 3 for our main model, optimal carbon pricing satisfies $\tau_{\text{co2}} = \mu/\Lambda$. In that case, $\mathrm{T} = \mathcal{T}'(y)$. Intuitively, the negative fiscal externality from lower carbon tax revenue exactly offsets the positive effect implied by the relaxation of the carbon cap constraint.

Under optimal carbon pricing, the gap between the EMTR and the welfare wedge is then given by:

$$\tau_{\text{eff}}(y) - \mathrm{T}(y) = \tau_{\text{co2}}\,(1-\mathcal{T}'(y))\, MPP(y).$$

This equation will be central for our results below. A larger gap implies a larger adjustment in the optimal degree of redistribution following an introduction of a binding exogenous limit on carbon emissions. Mechanically, the gap is increasing in (i) the level of the carbon tax and (ii) the marginal propensity to pollute. It decreases in the amount of redistribution as captured by the marginal income tax rate $\mathcal{T}'(y)$.

3 Model

3.1 Model Basics

Heterogeneity. We consider a static economy with a continuum of heterogeneous households of total mass one, indexed by their productivity $\theta \in [\underline{\theta}, \bar{\theta}]$. Let $H(\theta)$ denote the cumulative distribution function of productivity, with density $h(\theta)$. Households earn a gross labor income given by $y = l\theta$, where l denotes labor supply.

Preferences. Households make a labor-leisure decision and allocate their income across a vector of consumption goods and services $\mathbf{c} = (c_1, c_2, \ldots, c_I)$. Preferences over consumption and labor supply are represented by

$$U(\mathbf{c}) - \frac{l^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}},$$

where $\varphi > 0$ denotes the Frisch elasticity of labor supply. The utility function is weakly separable in consumption and leisure. In the absence of externalities, this implies that uniform commodity taxation is optimal (Atkinson and Stiglitz 1976).

In general, we allow for non-homothetic consumption preferences, which are necessary to produce empirically plausible carbon Engel curves.[6] For parts of the theoretical analysis, we impose the following functional form assumption on preferences, which sharpens the theoretical analysis and isolates the main mechanism.

Assumption 1. *Consumption utility is given by a homothetic CES aggregator,*

$$U(\mathbf{c}) = \left(\sum_{i=1}^{I} \Omega_i^{\frac{1}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}},$$

with $\sigma > 0$ and weights $\Omega_i > 0$.

Assumption 1 implies homothetic demands and, in particular, an indirect utility from consumption that is linear in total expenditure. Hence, the marginal utility of expenditure does not vary with expenditure implying zero income effects in labor supply. In the theoretical analysis, Assumption 1 serves two purposes. First, for the results on the equity-efficiency trade-off (Lemma 2 and Proposition 2), it is imposed only for expositional clarity; the corresponding results without this restriction are provided in Appendix A. Second, the sharp comparative statics results in Section 5.3 rely on this assumption.

Government. The government levies a nonlinear income tax schedule $\mathcal{T}(y)$ and a carbon tax τ_{co2} per unit of emission.[7]

Carbon emissions. Absent abatement, consumption of one unit of good i generates $\bar{f}_i > 0$ units of carbon emissions, where the emission intensity varies across goods. We then allow firms to reduce these baseline emissions through abatement, modeled in reduced form following the abatement block of the DICE model (Barrage and Nordhaus 2024). Firms choose abatement

[6] For recent papers that have studied Mirrleesian income taxation in the presence of non-homothetic preferences, see Jaravel and Olivi (2025) and Ferriere, Grübener, and Sachs (2025).

[7] Our results can easily be generalized to a setting with a uniform commodity tax as we show in an earlier version of this paper (Fricke, Fuest, and Sachs 2025).

$a_i \in [0, \bar{f}_i]$ at convex resource cost $\kappa_i(a_i)$ per unit. The resulting endogenous carbon emissions intensity for one unit of good i is $f_i = \bar{f}_i - a_i$.

Let q_i denote the marginal resource cost of producing one unit of good i before abatement costs and carbon taxes. Competitive firms take the carbon tax τ_{co2} as given and choose abatement to minimize the unit cost of good i. The resulting unit price is therefore

$$p_i(\tau_{\text{co}_2}) \equiv q_i + \min_{a_i \in [0, \bar{f}_i]} \left\{ \kappa_i(a_i) + \tau_{\text{co}_2}(\bar{f}_i - a_i) \right\},$$

and equals the marginal cost of production given optimal abatement. The optimal abatement $a_i^*(\tau_{\text{co}_2})$ is implicitly defined by $\kappa_i'(a_i^*(\tau_{\text{co}_2})) = \tau_{\text{co}_2}$: firms abate up to the point where the marginal cost of abatement equals the marginal savings in carbon tax payments. The resulting endogenous emissions intensity per unit of good i is then given by

$$f_i(\tau_{\text{co}_2}) \equiv \bar{f}_i - a_i^*(\tau_{\text{co}_2}).$$

We adopt a consumption-based accounting perspective: the residual emissions generated along the production chain of one unit of good i are attributed to the final consumption of that unit. Hence, after firms optimally choose abatement, household consumption of one unit of good i is associated with $f_i(\tau_{\text{co}_2})$ units of emissions, even though these emissions physically arise during production. Under perfect competition, the carbon tax is fully passed through to consumer prices, so purchasing one unit costs $p_i(\tau_{\text{co}_2})$.

By the envelope theorem, the effect of the carbon tax on the consumer price is given by

$$p_i'(\tau_{\text{co}_2}) = f_i(\tau_{\text{co}_2}). \tag{3}$$

That is, a marginal increase in the carbon tax raises the price of good i by exactly its residual emissions intensity. This identity will be useful below when characterizing how carbon taxes affect household carbon footprints and labor supply incentives.

Simplifying notation. Throughout, we hold fixed and suppress the production primitives $\{q_i, \bar{f}_i, \kappa_i\}_{i=1}^{I}$. Thus, any dependence on τ_{co_2} below incorporates the induced responses of both consumer prices $p_i(\tau_{\text{co}_2})$ and emission intensities $f_i(\tau_{\text{co}_2})$.

3.2 Household decision problem given taxes

We divide the household's decision problem into two parts. First, households choose their optimal consumption allocation given net income $e = y - \mathcal{T}(y)$. Second, households choose labor supply given the optimal consumption allocation.

The first stage is given by

$$u(e;\tau_{\text{co}_2}) = \max_{\mathbf{c}} U(\mathbf{c}) \quad \text{s.t.} \quad \sum_{i=1}^{I} c_i\, p_i(\tau_{\text{co}_2}) = e. \tag{4}$$

Since households differ only in productivity and consumption is weakly separable from labor, the consumption allocation problem for a given expenditure level e does not otherwise depend on θ. The first-order condition is:

$$\forall i: \ \lambda(e;\tau_{\text{co}_2}) = \frac{U_i\left(\mathbf{c}\left(e;\tau_{\text{co}_2}\right)\right)}{p_i(\tau_{\text{co}_2})}, \tag{5}$$

where $\lambda(e;\tau_{\text{co}_2})$ is the Lagrangian multiplier on the household's budget constraint. By the envelope theorem, it follows that $u_e(e;\tau_{\text{co}_2}) = \lambda(e;\tau_{\text{co}_2})$. Let $c_i(e;\tau_{\text{co}_2})$ denote the optimal consumption of good i, and define the resulting carbon footprint as

$$f(e;\tau_{\text{co}_2}) \equiv \sum_{i=1}^{I} f_i(\tau_{\text{co}_2})\, c_i(e;\tau_{\text{co}_2}).$$

Let $\mathcal{P} \equiv (\mathcal{T}, \tau_{\text{co}_2})$ denote the policy vector. The second stage of the household's problem is given by

$$V(\theta;\mathcal{P}) = \max_{y} u(e;\tau_{\text{co}_2}) - \frac{(y/\theta)^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}} \quad \text{s.t.} \quad e = y - \mathcal{T}(y). \tag{6}$$

The corresponding first-order condition for labor supply is given by

$$u_e(e(\theta,\mathcal{P});\tau_{\text{co}_2})\,\theta\,[1-\mathcal{T}'(y(\theta,\mathcal{P}))] = \left(\frac{y(\theta,\mathcal{P})}{\theta}\right)^{\frac{1}{\varphi}}. \tag{7}$$

In the following, we often use $u_e(\theta,\mathcal{P})$ as short-hand notation for $u_e(e(\theta,\mathcal{P});\tau_{\text{co}_2})$, and similarly we write $f(\theta,\mathcal{P}) \equiv f(e(\theta,\mathcal{P});\tau_{\text{co}_2})$.

Comparative statics. We now define the parameters that govern how households adjust their carbon emissions in response to a change in the carbon tax. First, we define the marginal propensity to pollute (MPP) as the partial effect of net income e on household emissions:[8]

$$MPP(\theta,\mathcal{P}) \equiv f_e(e(\theta,\mathcal{P});\tau_{\text{co}_2}).$$

Second, we define the partial effect of the carbon tax τ_{co_2} on household emissions, holding expenditures fixed, as

$$f_{\tau_{\text{co}_2}}(\theta,\mathcal{P}) \equiv \left.\frac{\partial f(e;\tau_{\text{co}_2})}{\partial \tau_{\text{co}_2}}\right|_{e=e(\theta,\mathcal{P})}.$$

[8]Under Assumption 1, demands are homothetic, so carbon footprints are proportional to expenditure. Hence $MPP(\theta,\mathcal{P})$ is constant across θ in that benchmark.

Note that $f_{\tau_{\mathrm{co}_2}}$ reflects both changes in the consumption allocation across goods and changes in emissions intensities through abatement. The total effect of a change in the carbon tax τ_{co_2} on household emissions is then given by:

$$\frac{df(\theta,\mathcal{P})}{d\tau_{\mathrm{co}_2}} = f_{\tau_{\mathrm{co}_2}}(\theta,\mathcal{P}) + MPP(\theta,\mathcal{P})\left(1-\mathcal{T}'(y(\theta,\mathcal{P}))\right)\frac{\partial y(\theta,\mathcal{P})}{\partial \tau_{\mathrm{co}_2}},$$

where the second term captures the indirect effect through the labor supply response to the carbon tax and the resulting change in the level of net income.

In the following, we derive an expression for the impact of carbon taxes on labor supply, $\frac{\partial y(\theta,\mathcal{P})}{\partial \tau_{\mathrm{co}_2}}$, which depends on classical labor supply elasticities. This helps to relate the labor supply effects of carbon taxes to the well-known labor supply effects of income taxes. First, the elasticity of income with respect to the retention rate $1-\mathcal{T}'$ and the marginal propensity to earn out of unearned income are given by

$$\varepsilon_{y,1-\mathcal{T}'}(\theta,\mathcal{P}) \equiv \frac{d\log(y(\theta,\mathcal{P}))}{d\log(1-\mathcal{T}')} \quad \text{and} \quad \eta(\theta,\mathcal{P}) \equiv \frac{\partial y(\theta,\mathcal{P})}{\partial \mathcal{T}(0)}.$$

The exact formal expressions for both objects are standard and can be found in Appendix A.1.1.

We now provide a result for the labor supply response to a carbon tax reform $d\tau_{\mathrm{co}_2}$:[9]

Lemma 1. *Consider a marginal reform around a given policy $\mathcal{P}$. Holding the income tax schedule fixed, a marginal increase in the carbon tax $d\tau_{\mathrm{co}_2}$ has the same first-order effects on utility and labor supply as the following income tax reform*

$$\forall\theta:\ d\mathcal{T}(y(\theta,\mathcal{P})) = f(\theta,\mathcal{P})d\tau_{\mathrm{co}_2}. \tag{8}$$

The change in income is given by:

$$\frac{\partial y(\theta,\mathcal{P})}{\partial \tau_{\mathrm{co}_2}}d\tau_{\mathrm{co}_2} = -\frac{\varepsilon_{y,1-\mathcal{T}'}(\theta,\mathcal{P})y(\theta,\mathcal{P})}{1-\mathcal{T}'(y(\theta,\mathcal{P}))}MPP(\theta,\mathcal{P})\left(1-\mathcal{T}'(y(\theta,\mathcal{P}))\right)d\tau_{\mathrm{co}_2} + \eta(\theta,\mathcal{P})f(\theta,\mathcal{P})d\tau_{\mathrm{co}_2}. \tag{9}$$

Proof. See Appendix A.1.2. □

Lemma 1 decomposes the labor supply effect of a carbon tax into two channels: (i) a substitution effect that operates through the marginal payment on additional earnings, and (ii) an income effect that operates through the reduction in disposable resources.

[9] This is equivalent to Lemma 1 of Saez (2002), who studies the desirability of non-uniform commodity taxes in the presence of preference heterogeneity.

The first term in (9) captures the substitution effect of the carbon tax on labor supply. A marginal increase in the carbon tax changes labor supply incentives in exactly the same way as an increase in the marginal income tax rate of[10]

$$d\mathcal{T}'(y(\theta,\mathcal{P})) = MPP(\theta,\mathcal{P})\,(1-\mathcal{T}'(y(\theta,\mathcal{P})))\,d\tau_{\mathrm{co}_2}. \tag{10}$$

To see this, note that one additional unit of gross income generates $(1-\mathcal{T}'(y(\theta,\mathcal{P})))$ units of extra expenditure. Since an extra unit of expenditure raises emissions by $MPP(\theta,\mathcal{P})$, the associated increase in carbon tax payments is $MPP(\theta,\mathcal{P})\,(1-\mathcal{T}'(y(\theta,\mathcal{P})))\,d\tau_{\mathrm{co}_2}$, which is exactly the same wedge created by the income tax reform in (10). This is analogous to the labor supply effects of savings taxation emphasized by Ferey, Lockwood, and Taubinsky (2024).

The second term in (9) captures the income effect. A marginal increase in the carbon tax reduces household income by $f(\theta,\mathcal{P})d\tau_{\mathrm{co}_2}$, generating an income effect equivalent to that of the income tax reform in (8).

3.3 Government and Policies

Welfare and constraints. The government maximizes a standard social welfare function,

$$\mathcal{W}(\mathcal{P}) = \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{G}\left(V(\theta,\mathcal{P})\right)\omega(\theta)\,dH(\theta), \tag{11}$$

where the weights $\omega(\theta) \geq 0$ are otherwise unrestricted and normalized to integrate to one. We assume that either (i) $\mathcal{G}(x) = x$ or (ii) $\mathcal{G}'(x) > 0$ and $\mathcal{G}''(x) < 0$. The latter captures aversion to utility inequality.

We define the social marginal utility of type θ as[11]

$$g(\theta,\mathcal{P}) = \frac{\partial \mathcal{W}(\mathcal{P})}{\partial e(\theta)}\frac{1}{h(\theta)} = \mathcal{G}'\left(V(\theta,\mathcal{P})\right)\omega(\theta)u_e(\theta,\mathcal{P}). \tag{12}$$

The government budget constraint is given by:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left(\mathcal{T}(y(\theta,\mathcal{P})) + \tau_{\mathrm{co}_2} f(\theta,\mathcal{P})\right)dH(\theta) \geq \mathcal{R}, \tag{13}$$

where $\mathcal{R}$ is an exogenous revenue requirement. The associated Lagrangian multiplier is denoted by Λ.

[10] Equation (10) follows directly from differentiating (8).

[11] This definition slightly differs from the standard definition of endogenous social welfare weights going back to Saez (2001), as we do not normalize them by the shadow value of public funds. In our optimal policy formulas below, the reader can interpret the social marginal utility $g(\theta,\mathcal{P})$ as being expressed in units of public funds, since all optimality conditions involve it only through relative terms.

Besides the government budget constraint, we assume that the government has to obey a carbon cap constraint

$$\int_{\underline{\theta}}^{\overline{\theta}} f(\theta, \mathcal{P}) dH(\theta) \equiv F \leq \bar{F}, \tag{14}$$

where $\bar{F}$ is an exogenous cap on aggregate carbon emissions. We denote the Lagrangian multiplier of the carbon cap constraint by μ. This constraint captures, for example, limits imposed by international carbon reduction agreements. All our results below would carry over if, instead of a cap, we modeled carbon emissions as an atmospheric externality, with welfare net of the externality given by $\mathcal{W} - D(F)$, where the damage function $D(F)$ is monotonically increasing in aggregate emissions F.

Effective marginal tax rate. We define the EMTR of household θ as the increase in total tax payment $\mathcal{T}(y(\theta, \mathcal{P})) + \tau_{\mathrm{co}_2} f(\theta, \mathcal{P})$ resulting from a one-unit increase in the household's gross labor income:

$$\tau_{\mathrm{eff}}(\theta, \mathcal{P}) = \mathcal{T}'(y(\theta, \mathcal{P})) + \left(1 - \mathcal{T}'(y(\theta, \mathcal{P}))\right) MPP(\theta, \mathcal{P})\, \tau_{\mathrm{co}_2}. \tag{15}$$

Besides the marginal income tax rate $\mathcal{T}'(\cdot)$, it also accounts for the additional tax revenue generated through higher carbon emissions resulting from a one-unit increase in gross household income.[12] Note that $\tau_{\mathrm{eff}}(\theta, \mathcal{P})$ is the relevant concept to capture the fiscal externalities that arise from changes in incomes of type θ.

3.4 Incidence of Carbon Taxes

We briefly discuss the incidence of a carbon tax. By Lemma 1, a carbon tax increase has the same first-order effects on utility and earnings as the income tax reform in (8). The remaining difference is that the carbon tax changes relative consumer prices and induces abatement. To isolate this difference, combine a marginal increase $d\tau_{\mathrm{co}_2}$ with the opposite of the equivalent income tax reform,

$$d\mathcal{T}(y(\theta, \mathcal{P})) = -f(\theta, \mathcal{P}) d\tau_{\mathrm{co}_2}.$$

This joint perturbation leaves utility and earnings unchanged to first order. At fixed earnings, the offsetting income tax reform increases expenditure by $f(\theta, \mathcal{P}) d\tau_{\mathrm{co}_2}$. The household's carbon footprint therefore changes by

$$df^c(\theta, \mathcal{P}) = \left[f_{\tau_{\mathrm{co}_2}}(\theta, \mathcal{P}) + MPP(\theta, \mathcal{P}) f(\theta, \mathcal{P})\right] d\tau_{\mathrm{co}_2} \equiv f^c_{\tau_{\mathrm{co}_2}}(\theta, \mathcal{P}) d\tau_{\mathrm{co}_2}.$$

[12] Under Assumption 1, $MPP(\theta, \mathcal{P})$ is constant across types. In that benchmark, cross-sectional variation in the carbon component of $\tau_{\mathrm{eff}}(\theta, \mathcal{P})$ arises only through the net-of-tax factor $1 - \mathcal{T}'(y(\theta, \mathcal{P}))$.

The object $f^c_{\tau_{\text{co}_2}}$ is the compensated carbon footprint response: it captures consumption reallocation and abatement while compensating the household for the first-order purchasing power loss. The associated fiscal externality is

$$\tau_{\text{co}_2} \int_{\underline{\theta}}^{\overline{\theta}} f^c_{\tau_{\text{co}_2}}(\theta, \mathcal{P})\, dH(\theta)\, d\tau_{\text{co}_2}. \tag{16}$$

With a slack carbon cap, the welfare effect of the joint perturbation is Λ times (16). Since compensated responses are negative, $f^c_{\tau_{\text{co}_2}}(\theta, \mathcal{P}) < 0$, any positive carbon tax can be improved upon by lowering it. Thus, the optimal policy satisfies $\tau_{\text{co}_2} = 0$. Put differently, when the carbon cap constraint (14) is slack, a positive carbon tax is strictly inferior to the distributionally equivalent income tax adjustment.[13]

In Section 4, we first consider this benchmark without a binding carbon cap and characterize the optimal 'status quo' policy $\mathcal{P}_{sq} = (\mathcal{T}_{sq}, 0)$. In Section 5, we introduce a binding carbon cap, which makes it optimal for the government to impose a positive carbon tax. We characterize the resulting optimal policy as $\mathcal{P}_{co_2} = (\mathcal{T}_{co_2}, \tau_{\text{co}_2})$.[14] We interpret the difference between the two income tax schedules as the optimal carbon tax rebate:

$$R_{co_2}(y) \equiv \mathcal{T}_{\text{sq}}(y) - \mathcal{T}_{co_2}(y),$$

and derive its key properties. We also examine whether $\mathcal{P}_{co_2}$ is more or less redistributive than the status quo policy $\mathcal{P}_{sq}$.

4 Normative Benchmark: No Carbon Cap

As a first step, we characterize the optimal tax-transfer system without a binding carbon cap, which we refer to as the status quo. As shown in Section 3.4, the optimal carbon tax is zero in this benchmark. The government therefore chooses only the nonlinear income tax schedule and solves

$$\max_{\mathcal{T}(\cdot)} \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{G}\left(V(\theta, \mathcal{P})\right) \omega(\theta) dH(\theta) \tag{17}$$

subject to the government budget constraint (13) and optimal household behavior (4) and (6), where $\mathcal{P} = (\mathcal{T}, 0)$ for every candidate schedule $\mathcal{T}$. We denote the solution by $\mathcal{P}_{\text{sq}} = (\mathcal{T}_{\text{sq}}, 0)$.

[13] This logic mirrors the Atkinson–Stiglitz result for differentiated commodity taxation: such taxes cannot achieve distributional goals or influence labor supply incentives beyond what the nonlinear income tax can achieve, but they introduce additional distortions in the consumption allocation (Atkinson and Stiglitz 1976).

[14] This involves a slight abuse of notation: for $\mathcal{P}$ and $\mathcal{T}$, the subscript $co2$ denotes optimality under the carbon cap constraint, whereas for the carbon tax τ_{co_2}, the subscript was used earlier even when the carbon tax was not set optimally.

The following lemma states the optimal income tax condition in the form presented by Heathcote and Tsujiyama (2021). We impose Assumption 1 in the main text to suppress income effects and keep the condition transparent. Appendix A.2 derives the corresponding formula with general income effects.

Lemma 2. *Assume that Assumption 1 holds. The optimal income tax schedule $\mathcal{T}_{\text{sq}}$ that solves (17) satisfies*

$$\forall\, \theta^* :\; D(\theta^*, \mathcal{P}_{\text{sq}}) = E(\theta^*, \mathcal{P}_{\text{sq}}),$$

where

$$D(\theta^*, \mathcal{P}_{\text{sq}}) = 1 - \frac{\mathbb{E}[g(\theta, \mathcal{P}_{\text{sq}})|\theta \geq \theta^*]}{\mathbb{E}[g(\theta, \mathcal{P}_{\text{sq}})]} \tag{18}$$

captures the distributional gains of raising the marginal tax rate $\mathcal{T}'(y(\theta))$ and

$$E(\theta^*, \mathcal{P}_{\text{sq}}) = \frac{\tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{sq}})}{1 - \mathcal{T}'_{\text{sq}}(y(\theta^*))} \frac{\varphi}{1+\varphi} \frac{h(\theta^*)\theta^*}{1 - H(\theta^*)} \tag{19}$$

captures the efficiency cost of raising the marginal tax rate.

Proof. See Appendix A.2. □

In this benchmark, $\tau_{\text{co2}} = 0$ and thus $\tau_{\text{eff}}(\theta, \mathcal{P}_{\text{sq}}) = \mathcal{T}'_{\text{sq}}(y(\theta))$. This highlights the usual benchmark: absent a binding carbon constraint, the EMTR governs both marginal redistribution and marginal efficiency costs. Concretely, the optimal marginal tax rate $\mathcal{T}'_{\text{sq}}(y(\theta^*))$ balances optimally the distributional gains $D(\theta^*, \mathcal{P}_{\text{sq}})$ with the efficiency cost $E(\theta^*, \mathcal{P}_{\text{sq}})$. The term $D(\theta^*, \mathcal{P}_{\text{sq}})$ measures the proportional gap between the population-average social marginal utility and the average social marginal utility of individuals earning more than $y(\theta^*, \mathcal{P}_{\text{sq}})$. Thus, it captures the gains of redistributing incomes from individuals with income above $y(\theta^*, \mathcal{P}_{\text{sq}})$ to those with income below $y(\theta^*, \mathcal{P}_{\text{sq}})$.

The term $E(\theta^*, \mathcal{P}_{\text{sq}})$ captures the welfare cost generated by the induced labor supply response. Since income changes affect welfare through fiscal externalities, the earnings response is weighted by $\tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{sq}})$. Under Assumption 1, income effects are zero, so (19) contains only the standard substitution effect.

Figure 1 schematically illustrates the trade-off for a given type θ. The efficiency costs (red curve) increase in $\tau_{\text{eff}}(\theta)$. The distributional gains term (black curve) decreases in $\tau_{\text{eff}}(\theta)$, which holds whenever the welfare weights $g(\theta, \mathcal{P}_{\text{sq}})$ are decreasing. Note that the distributional gains term depends on the lump sum element, the EMTRs at all other levels and the resulting allocation. The illustrative graph will be useful for the interpretation of our main comparative statics results in Section 5.3.

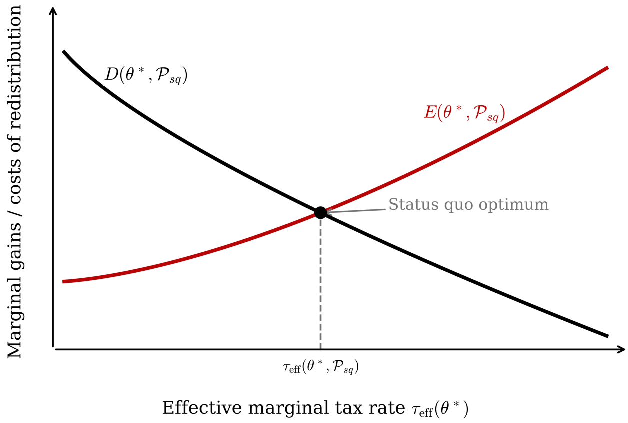

Figure 1: Status-quo equity–efficiency trade-off

Notes: In the absence of a binding carbon cap, marginal redistribution is governed by the standard Mirrleesian condition $D(\theta^*, \mathcal{P}_{sq}) = E(\theta^*, \mathcal{P}_{sq})$. The black curve shows the marginal gains from redistribution, the red curve the corresponding marginal efficiency costs. Their intersection determines the status-quo optimal tax system and the associated EMTR $\tau_{\text{eff}}(\theta^*, \mathcal{P}_{sq})$.

5 Optimal Carbon Rebate Policies

We now turn to the case in which the government faces a binding carbon constraint (14). The optimal policy consists of a nonlinear income tax schedule $\mathcal{T}(y)$ and a carbon tax τ_{co2} that together ensure compliance with the emission cap. The optimal implementation of the carbon cap requires solving the following problem:

$$\max_{\mathcal{T}(\cdot), \tau_{\text{co2}}} \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{G}\left(V(\theta; \mathcal{P})\right) \omega(\theta) \, dH(\theta) \tag{20}$$

subject to the carbon cap constraint (14), the government budget constraint (13), and optimal household behavior (4) and (6). As opposed to (17), the government now optimally chooses both elements of the policy vector $\mathcal{P} = (\mathcal{T}, \tau_{\text{co2}})$. We denote the solution to this problem by $\mathcal{P}_{\text{co2}} = (\mathcal{T}_{\text{co2}}, \tau_{\text{co2}})$.

Note that the differences between $\mathcal{T}_{\text{co2}}$ and $\mathcal{T}_{\text{sq}}$ arise from the carbon cap constraint and the associated use of the carbon tax. Equivalently, we can fix the income tax schedule at $\mathcal{T}_{\text{sq}}$ and introduce a carbon tax rebate function $R_{\text{co2}}(y)$ defined by

$$\forall y: \; R_{\text{co2}}(y) = \mathcal{T}_{\text{sq}}(y) - \mathcal{T}_{\text{co2}}(y).$$

This representation facilitates the analysis of how carbon tax revenues should be optimally rebated across the income distribution.

We present our results sequentially. Section 5.1 analyzes the optimal carbon tax τ_{co2}. Section 5.2 examines the equity-efficiency trade-off. Finally, Section 5.3 explores the properties of the carbon rebate function and assesses whether the optimal policy $\mathcal{P}_{\text{co2}}$ is more redistributive than $\mathcal{P}_{\text{sq}}$.

5.1 Optimal carbon tax

Lemma 3. *The carbon tax rate that solves the optimization problem (20) is given by*

$$\tau_{\text{co}_2} = \frac{\mu}{\Lambda}, \tag{21}$$

where Λ is the Lagrangian multiplier on the government budget constraint (13) and μ is the Lagrangian multiplier on the carbon cap constraint (14).

Proof. See Appendix A.3.1. □

This formula for the carbon tax represents a first-best rule: it holds even if the government could implement type-dependent lump-sum taxes. This result resembles Kaplow (2012) and Jacobs and De Mooij (2015), who show that – under the Atkinson-Stiglitz assumptions – the optimal Pigouvian tax follows a first-best rule, even in a second-best setting with income taxation.[15]

To see the logic behind Lemma 3, consider a marginal increase in the carbon tax $d\tau_{\text{co2}}$ combined with the opposite of the equivalent income tax reform in (8),

$$d\mathcal{T}\big(y(\theta, \mathcal{P}_{\text{co2}})\big) = -f(\theta, \mathcal{P}_{\text{co2}})\, d\tau_{\text{co2}}.$$

By Lemma 1, this joint perturbation leaves utility and earnings unchanged to first order. At fixed earnings, the income tax adjustment raises expenditure by $f(\theta, \mathcal{P}_{\text{co2}})d\tau_{\text{co2}}$. Hence, the remaining change in the household's carbon footprint is the compensated response introduced in Section 3.4,

$$df(\theta, \mathcal{P}_{\text{co2}}) = f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P}_{\text{co2}})d\tau_{\text{co2}}.$$

The direct increase $f(\theta, \mathcal{P}_{\text{co2}})d\tau_{\text{co2}}$ in carbon tax payments is exactly offset by the income tax adjustment. The remaining budget effect is therefore $\tau_{\text{co2}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P}_{\text{co2}})d\tau_{\text{co2}}$, while the corresponding change in the carbon cap is $f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P}_{\text{co2}})d\tau_{\text{co2}}$. The first-order condition is thus

$$(\Lambda\tau_{\text{co2}} - \mu) \int_{\underline{\theta}}^{\overline{\theta}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P}_{\text{co2}})\, dH(\theta)\, d\tau_{\text{co2}} = 0. \tag{22}$$

Since the aggregate compensated carbon footprint response is nonzero, equation (22) yields (21).

[15] As shown by Jacobs and De Mooij (2015), this no longer holds when the utility function is non-separable.

5.2 Equity-Efficiency Trade-Off

We now derive the analogue of Lemma 2 under a binding emissions cap. The key difference will be that labor supply responses affect welfare not only through tax revenue, but also through their impact on the tightness of the emissions constraint. We first formalize this in Proposition 1 where we characterize the welfare effect of a marginal increase in gross labor income and compare it to the fiscal externality.

Proposition 1. *At the optimal carbon constrained policy $\mathcal{P}_{\text{co2}}$, a marginal policy induced change in gross labor income $dy(\theta)$ generates a fiscal externality and a carbon cap externality. Its first-order effect on the government's Lagrangian is*

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial y(\theta)} dy(\theta) &= \Lambda \left[\underbrace{\tau_{\text{eff}}(\theta, \mathcal{P}_{\text{co2}})}_{\textit{fiscal externality}} - \underbrace{\frac{\mu}{\Lambda} \left(1 - \mathcal{T}'_{\text{co2}}\big(y(\theta, \mathcal{P}_{\text{co2}})\big)\right) MPP(\theta, \mathcal{P}_{\text{co2}})}_{\textit{carbon cap externality}} \right] dy(\theta) \\ &\equiv \Lambda \qquad \underbrace{\mathrm{T}(\theta, \mathcal{P}_{\text{co2}})}_{\textit{welfare-relevant labor wedge}} \qquad dy(\theta). \end{aligned} \tag{23}$$

Using Lemma 3, the welfare-relevant labor wedge simplifies to

$$\begin{aligned}\mathrm{T}(\theta, \mathcal{P}_{\text{co2}}) &= \tau_{\text{eff}}(\theta, \mathcal{P}_{\text{co2}}) - \tau_{\text{co2}} \left(1 - \mathcal{T}'_{\text{co2}}\big(y(\theta, \mathcal{P}_{\text{co2}})\big)\right) MPP(\theta, \mathcal{P}_{\text{co2}}) \\ &= \mathcal{T}'_{\text{co2}}\big(y(\theta, \mathcal{P}_{\text{co2}})\big). \end{aligned} \tag{24}$$

Consequently, whenever $\tau_{\text{co2}} > 0$, $MPP(\theta, \mathcal{P}_{\text{co2}}) > 0$, and $1 - \mathcal{T}'_{\text{co2}}(y(\theta, \mathcal{P}_{\text{co2}})) > 0$, the welfare-relevant labor wedge is strictly smaller than the EMTR.

Proof. By the envelope theorem, the induced earnings change has no direct first-order effect on household utility. A one-unit increase in gross earnings raises government revenue by $\tau_{\text{eff}}(\theta, \mathcal{P}_{\text{co2}})$ and raises emissions by

$$\left(1 - \mathcal{T}'_{\text{co2}}\big(y(\theta, \mathcal{P}_{\text{co2}})\big)\right) MPP(\theta, \mathcal{P}_{\text{co2}}).$$

Weighting these two effects by Λ and μ, respectively, gives (23). Equation (24) follows from Lemma 3 and the definition of the effective marginal tax rate in (15). $\square$

We call the gap between the EMTR and the welfare-relevant labor wedge the *carbon redistribution wedge* and denote it by

$$\begin{aligned}\Delta_{\text{red}}(\theta, \mathcal{P}_{\text{co2}}) &\equiv \tau_{\text{eff}}(\theta, \mathcal{P}_{\text{co2}}) - \mathrm{T}(\theta, \mathcal{P}_{\text{co2}}) \\ &= \tau_{\text{co2}} \left[1 - \mathcal{T}'_{\text{co2}}(y(\theta, \mathcal{P}_{\text{co2}}))\right] MPP(\theta, \mathcal{P}_{\text{co2}}). \end{aligned} \tag{25}$$

Thus, the EMTR exceeds the welfare-relevant labor wedge by the marginal carbon tax payment generated by an additional unit of gross earnings. A reduction in labor supply lowers both tax revenue and emissions. Under optimal carbon pricing, the welfare gain from relaxing the carbon cap exactly offsets the loss of carbon tax revenue, leaving only the marginal income tax rate as the welfare-relevant labor wedge. The object Δ_{red} measures the resulting reduction in the welfare cost of labor supply responses for a given EMTR.

We are now ready to present the counterpart to Lemma 2 under a binding emissions cap. As for Lemma 2, we consider no income effects in the main text by imposing Assumption 1. In Appendix A.3.2, we present the more general results.

Proposition 2. *Assume that Assumption 1 holds. Consider a government that solves (20). The optimal carbon tax satisfies (21). The optimal income tax schedule $\mathcal{T}_{\text{co2}}$ satisfies*

$$\forall\, \theta^* :\ D(\theta^*, \mathcal{P}_{\text{co2}}) = E(\theta^*, \mathcal{P}_{\text{co2}}),$$

where

$$D(\theta^*, \mathcal{P}_{\text{co2}}) = 1 - \frac{\mathbb{E}[g(\theta, \mathcal{P}_{\text{co2}})|\theta \geq \theta^*]}{\mathbb{E}[g(\theta, \mathcal{P}_{\text{co2}})]} \tag{26}$$

and

$$E(\theta^*, \mathcal{P}_{\text{co2}}) = \frac{\tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{co2}}) - \Delta_{\text{red}}(\theta^*, \mathcal{P}_{\text{co2}})}{1 - \mathcal{T}'_{\text{co2}}(y(\theta^*, \mathcal{P}_{\text{co2}}))} \frac{\varphi}{1+\varphi} \frac{h(\theta^*)\theta^*}{1 - H(\theta^*)} \tag{27}$$

captures the associated marginal efficiency cost.

Proof. See Appendix A.3.2. □

The key difference to Lemma 2 lies in the efficiency cost term $E_{\text{co2}}(\theta^*, \mathcal{P}_{\text{co2}})$. Labor supply responses are weighted by $\mathrm{T}(\theta^*, \mathcal{P}_{\text{co2}}) = \tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{co2}}) - \Delta_{\text{red}}(\theta^*, \mathcal{P}_{\text{co2}})$ rather than by $\tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{co2}})$. Hence, for a given allocation and $\tau_{\text{eff}}(\theta^*, \mathcal{P}_{\text{co2}})$, a positive carbon redistribution wedge lowers the marginal efficiency cost of redistribution. By contrast, the functional form of the distributional gains term is unchanged, although its value may differ because the carbon policy changes the allocation and social marginal utilities.[16] Hence, Proposition 2 shows that the carbon tax τ_{co2} asymmetrically affects the equity-efficiency trade-off: for a given allocation, it reduces the efficiency costs of redistribution while leaving the welfare gains from redistribution unchanged.

This distinction can be obscured by the fact that optimal carbon pricing implies $\mathrm{T}(\theta, \mathcal{P}_{\text{co2}}) = \mathcal{T}'_{\text{co2}}(y(\theta, \mathcal{P}_{\text{co2}}))$. Accordingly, the optimal income tax condition retains the standard Mirrlees–

[16] The ratio of social marginal utilities can change through two conceptually distinct channels: First, as in standard redistributive tax models, carbon tax and rebate policies change households' disposable resources and utility levels and thereby affect both $G'(V)$ and u_e. Second, if carbon Engel curves are nonlinear, a carbon tax generates heterogeneous inflation rates across the income distribution, which additionally affect the relative values of u_e. As shown by Jaravel and Olivi (2025), heterogeneous inflation rates have subtle implications for optimal redistributive policy because the marginal value of one dollar, measured in terms of consumption baskets, changes differently across the income distribution. By assuming linear carbon Engel curves for most of the theoretical analysis in the main text, we abstract from this channel. In the quantitative section below, the channel is present, however.

Saez form when written in terms of the marginal income tax rate.[17] This could mistakenly suggest that optimal redistribution and optimal carbon pricing are tangential.

Figure 1 provides a useful graphical interpretation of Proposition 2. Consider a given type θ and hold the allocation fixed. The marginal efficiency cost is governed by the welfare-relevant labor wedge $\mathrm{T}(\theta, \mathcal{P}_{\mathrm{co2}})$, whereas the associated EMTR is higher by $\Delta_{\mathrm{red}}(\theta, \mathcal{P}_{\mathrm{co2}})$:

$$\tau_{\mathrm{eff}}(\theta, \mathcal{P}_{\mathrm{co2}}) = \mathrm{T}(\theta, \mathcal{P}_{\mathrm{co2}}) + \Delta_{\mathrm{red}}(\theta, \mathcal{P}_{\mathrm{co2}}).$$

Hence, in Figure 1, the efficiency cost curve is shifted to the right by the carbon redistribution wedge. The same marginal efficiency cost is therefore consistent with a higher degree of effective redistribution. This graphical argument captures the basic force toward more redistribution. Clean comparative statics are nevertheless not immediate, because the size of the carbon redistribution wedge, the allocation, and hence both the distributional gains and efficiency cost terms adjust endogenously with the tax schedule across the income distribution. In the next subsection, we impose additional structure to derive sharp comparative statics for the change in redistribution and for the optimal carbon rebate schedule.

5.3 Comparative Statics

Building on Proposition 2, the preceding subsection showed that, for a given allocation and a given EMTR, the carbon redistribution wedge lowers the welfare-relevant efficiency cost of redistribution. We now use the optimality conditions in Lemma 2 and Proposition 2 to evaluate the direction of welfare improving reforms at two benchmark implementations of a marginal carbon cap.

We address two broader questions:

1. Should redistribution decrease or increase when a carbon cap is introduced?

2. How should carbon tax revenues be rebated across the income distribution?

Our propositions provide local answers to these broader questions by evaluating the optimality condition $D(\theta, \mathcal{P}) = E(\theta, \mathcal{P})$ from Lemma 2 and Proposition 2 at two reference implementations of the carbon cap. For the first question, the distributionally neutral implementation serves as a reference point rather than as a policy proposal. Because it leaves EMTRs at their status quo levels, it allows us to isolate whether the carbon cap itself creates a welfare-improving force toward more redistribution.

[17]The income tax condition in Proposition 4 of Jacobs and de Mooij (2015, equation (38)) is a close counterpart to our Proposition 2. They show that the optimal nonlinear income tax formula retains the standard Mirrlees–Saez form when expressed in terms of the marginal income tax rate. This corresponds to our result that, under optimal carbon pricing, the welfare-relevant labor wedge equals the marginal income tax rate. The additional implication in our analysis arises from distinguishing this wedge from the total EMTR, which also includes the marginal carbon tax burden. This distinction permits us to compare effective redistribution across the carbon constrained and unconstrained optima and to derive the optimal rebate schedule.

For the second question, the carbon dividend serves as the natural benchmark for rebate design. Because it returns the same amount to every household and therefore leaves the income gradient in carbon-tax liabilities fully reflected in net burdens, it allows us to ask whether the rebate should instead rise with income and partially offset this gradient.

To obtain tractable comparative statics, we impose two more simplifying assumptions. These assumptions apply only to the theoretical results in this section and are *not* imposed in the quantitative analysis in Section 6.

Assumption 2. *The carbon emission cap $\bar{F}$ is marginally binding: $\bar{F} = F^{\text{sq}} - \varepsilon$, where $\varepsilon \to 0$ and F^{sq} denotes the aggregate carbon footprint under the status quo policy $\mathcal{P}_{\text{sq}}$.*

This assumption allows us to focus on the first-order effects of carbon taxes on individual utility and behavior as well as government revenue. Accordingly, Proposition 3 and Proposition 5 characterize first-order welfare gradients at the respective benchmark policies. Proposition 4, by contrast, does not require Assumption 2.

Assumption 3. *The primitives of the economy and the welfare function are such that the optimal status quo income tax schedule $\mathcal{T}_{\text{sq}}$ is affine, i.e. it features a constant marginal tax rate: $\mathcal{T}'_{\text{sq}}(y) = \mathcal{T}'_{\text{sq}}(y')$ for all y, y' and a lump-sum transfer $b > 0$.*

This assumption does *not* restrict the structure of $\mathcal{T}'_{\text{co2}}$; it merely implies that the social planner optimally chooses an affine income tax schedule in the absence of a carbon cap. It is helpful for deriving sharper results, as it ensures that a household's marginal income tax rate $\mathcal{T}'_{\text{sq}}(y)$ is unaffected by small income changes induced by the carbon tax and the corresponding adjustments in the lump-sum rebate.

5.3.1 Increasing Redistribution is Welfare Improving

We first evaluate the optimality condition at a distributionally neutral benchmark. Let $\mathcal{P}_{\text{dn}} = (\mathcal{T}_{\text{dn}}, \tau_{\text{co2}})$ denote a candidate policy that implements the marginal carbon cap and satisfies the carbon tax condition in Lemma 3, and let $R_{\text{dn}}(y) \equiv \mathcal{T}_{\text{sq}}(y) - \mathcal{T}_{\text{dn}}(y)$ denote the corresponding rebate. The benchmark satisfies

$$\forall\, \theta : \; R_{\text{dn}}(y(\theta, \mathcal{P}_{\text{dn}})) = \tau_{\text{co2}}\, f(\theta, \mathcal{P}_{\text{dn}}). \tag{28}$$

To first order, the rebate exactly offsets each household's carbon tax burden. Utility and earnings therefore remain at their status quo levels to first order, while Assumption 3 implies that EMTRs are unchanged:

$$\forall\, \theta : \; \tau_{\text{eff}}(\theta, \mathcal{P}_{\text{dn}}) = \tau_{\text{eff}}(\theta, \mathcal{P}_{\text{sq}})$$

We now evaluate the optimality condition in Proposition 2 at this benchmark.

Proposition 3. *Assume that Assumptions 1, 2, and 3 hold. At a distributionally neutral implementation $\mathcal{P}_{\text{dn}}$ satisfying (28),*

$$E(\theta, \mathcal{P}_{\text{dn}}) < D(\theta, \mathcal{P}_{\text{dn}}) \qquad \forall \theta \in]\underline{\theta}, \overline{\theta}[.$$

Hence, $\mathcal{P}_{\text{dn}}$ is not locally optimal. Starting from this benchmark, a marginal reform that decreases $R'_{\text{dn}}(y(\theta))$ – equivalently, increases $\mathcal{T}'_{\text{dn}}(y(\theta))$ and the EMTR – raises welfare.

Proof. See Appendix A.3.3. □

We first consider the distributional gains term. Since the distributionally neutral policy leaves utility levels unchanged to first order, only the change in the marginal utility of expenditure $u_e(\theta)$ can matter. As shown in Appendix A.3.3,

$$\forall \theta : \quad u_e(\theta, \mathcal{P}_{\text{dn}}) = u_e(\theta, \mathcal{P}_{\text{sq}}) \left(1 - \tau_{\text{co2}} \, MPP(\mathcal{P}_{\text{sq}})\right). \tag{29}$$

Here $MPP(\mathcal{P}_{\text{sq}}) \equiv MPP(\theta, \mathcal{P}_{\text{sq}})$ is constant across θ under Assumption 1. Equation (29) scales all social marginal utilities by the same factor to first order and therefore leaves their relative values unchanged. Hence,

$$D(\theta, \mathcal{P}_{\text{dn}}) = D(\theta, \mathcal{P}_{\text{sq}}) \qquad \forall \theta.$$

Next, we turn to the efficiency cost term. Distributional neutrality keeps the EMTRs at their status quo level, but Proposition 2 implies that the welfare-relevant labor wedge is lower than the EMTR by the carbon redistribution wedge. Equivalently, equality of the EMTRs requires

$$\mathcal{T}'_{\text{dn}}(y(\theta)) = \mathcal{T}'_{\text{sq}}(y(\theta)) - (1 - \mathcal{T}'_{\text{dn}}(y(\theta))) \, MPP(\theta, \mathcal{P}_{\text{dn}}) \tau_{\text{co2}}. \tag{30}$$

Substituting (30) into $\mathrm{T}(\theta, \mathcal{P}_{\text{dn}})$ and (27) yields

$$E(\theta, \mathcal{P}_{\text{dn}}) < E(\theta, \mathcal{P}_{\text{sq}}).$$

Combining this inequality with the status quo condition in Lemma 2 and the unchanged distributional gains term gives $E(\theta, \mathcal{P}_{\text{dn}}) < D(\theta, \mathcal{P}_{\text{dn}})$ to first-order. Thus, a marginal increase in the EMTR raises welfare.

Figure 2(a) illustrates this argument using the equity–efficiency diagram from the preceding subsection. At the distributionally neutral benchmark, the EMTR remains at its status quo value. The carbon redistribution wedge shifts the efficiency cost curve to the right. At the unchanged EMTR, marginal efficiency costs are therefore below marginal distributional gains, so a marginal increase in the EMTR raises welfare. The intersection of the shifted efficiency cost curve and the distributional gains curve is only schematic. The diagram is drawn for a given income level and holds the rest of the tax schedule fixed. Moving to the right at that point changes the allocation and social marginal utilities throughout the income distribution, which in turn shifts the distributional gains curves relevant at other income levels.

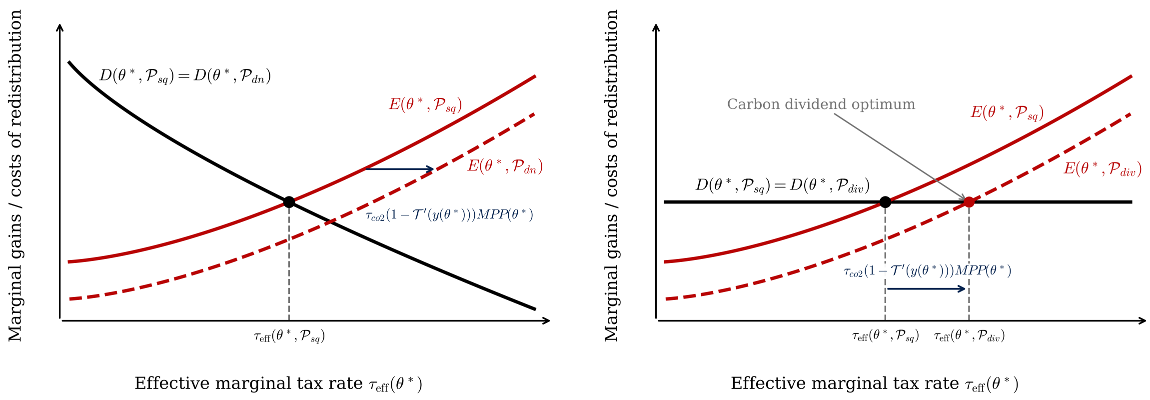

(a) Distributionally neutral rebate (b) Carbon dividend

Figure 2: Equity–efficiency trade-off in case of carbon cap

Notes: Panel (a) illustrates Proposition 3, and panel (b) illustrates Proposition 4. The black dot illustrates the status-quo optimum. The dashed red curves illustrate the lowered efficiency-cost curve. The red dot in the right panel illustrates the carbon dividend optimum.

5.3.2 The Carbon Dividend Benchmark

We next consider the opposite benchmark. Let $\mathcal{P}_{\text{div}} = (\mathcal{T}_{\text{div}}, \tau_{\text{co2}})$ denote a candidate policy that implements the carbon cap, satisfies the carbon tax condition in Lemma 3, and rebates carbon tax revenue through a budget-neutral carbon dividend,

$$R_{\text{div}}(y) \equiv \mathcal{T}_{\text{sq}}(y) - \mathcal{T}_{\text{div}}(y) = \bar{R}_{\text{div}}, \qquad R'_{\text{div}}(y) = 0.$$

We first identify the polar case in which this benchmark is exactly optimal and then evaluate the local direction of reform when marginal distributional gains decline.

Proposition 4. *Assume that Assumptions 1 and 3 hold and that $\mathcal{G}(V) = V$. Then, the optimal rebate is lump-sum; that is, a carbon dividend is optimal.*

Under Assumption 1, indirect utility is linear in expenditure, so u_e is common across types. With $\mathcal{G}(V) = V$, this common factor cancels from the distributional gains term, which is therefore policy invariant. Hence,

$$D(\theta, \mathcal{P}_{\text{div}}) = D(\theta, \mathcal{P}_{\text{sq}}) = E(\theta, \mathcal{P}_{\text{sq}}) = E(\theta, \mathcal{P}_{\text{div}}) \qquad \forall \theta.$$

The first equality follows from the policy invariance of the distributional gains term, and the second by construction from the optimality of the status-quo policy. For the final equality, note that a carbon dividend changes the income tax schedule only through its lump-sum component, so $\mathcal{T}'_{\text{div}}(y) = \mathcal{T}'_{\text{sq}}(y)$ at any given income level. Assumption 3 further implies that this marginal tax rate is constant in income. Hence, although the carbon tax lowers earnings, households continue to face the same marginal income tax rate as under the status quo.

The result is illustrated in Figure 2(b). Under $\mathcal{G}(V) = V$, the marginal gains from redistribution are constant: the distributional gains curve is horizontal and does not shift when the level of redistribution changes. In particular, changes in marginal tax rates at other income levels do not affect the position of this curve: under Assumption 1 and $\mathcal{G}(V) = V$, relative social marginal utilities are policy invariant. Hence, once the carbon cap lowers the welfare-relevant efficiency cost of redistribution, there is no countervailing decline in marginal distributional gains that would limit the redistributive adjustment.

When $\mathcal{G}''(V) < 0$, the distributional gains curve is decreasing in the EMTR and a carbon dividend is no longer optimal. A carbon dividend increases redistribution relative to the status quo. With diminishing social marginal utility, this changes the marginal distributional gains from further redistribution. Hence, the carbon dividend overshoots the redistributive adjustment that is optimal under the carbon cap. We show this formally by considering a budget-neutral carbon dividend and verifying that it violates the optimality condition in Proposition 2.

Proposition 5. *Assume that Assumptions 1, 2, and 3 hold and that $\mathcal{G}(V) = \frac{V^{1-\gamma_G}}{1-\gamma_G}$ with $\gamma_G > 0$ and with $\mathcal{G}(V) = \log V$ for $\gamma_G = 1$. The optimal policy $\mathcal{P}_{\text{co2}} = (\mathcal{T}_{\text{co2}}, \tau_{\text{co2}})$ cannot be consistent with a lump-sum rebate. At a budget-neutral carbon dividend $\mathcal{P}_{\text{div}}$,*

$$E(\theta, \mathcal{P}_{\text{div}}) > D(\theta, \mathcal{P}_{\text{div}}) \qquad \forall \theta \in]\underline{\theta}, \overline{\theta}[.$$

Hence, starting from a budget-neutral carbon dividend, the government would have an incentive to reduce redistribution. Equivalently, it would increase $R'_{\text{co2}}(y(\theta))$, or decrease $\mathcal{T}'_{\text{co2}}(y(\theta))$.

Proof. See Appendix A.3.4. □

Graphically, Proposition 5 can be read as a departure from the polar case in Figure 2(b). In Figure 2(b), the carbon dividend is optimal because the distributional gains curve is horizontal. When $\mathcal{G}''(V) < 0$, the carbon dividend reduces marginal distributional gains. Hence, at the carbon dividend point in Figure 2(b), the distributional gains curve would lie below the horizontal benchmark, while the marginal efficiency cost is unchanged. The planner therefore moves left from the carbon dividend allocation, choosing less redistribution than under a lump-sum rebate.[18]

[18] If relative utility-inequality aversion decreases with utility, the direction of the local reform can be reversed only at sufficiently high incomes, where households are net contributors under the carbon dividend. In that case, higher EMTRs and a locally decreasing rebate may be welfare improving. The broader conclusion that the carbon dividend entails too much redistribution over at least part of the income distribution, and is therefore not locally optimal, remains unaffected. See Appendix A.3.4 for details.

Taken together, Proposition 3 and Proposition 5 characterize the direction of welfare improving reforms at the two benchmark policies. At these benchmarks, the rebate slopes are

$$R'_{\text{div}}(y) = 0, \qquad R'_{\text{dn}}(y(\theta)) = \tau_{\text{co2}}\big(1 - \mathcal{T}'_{\text{dn}}(y(\theta))\big)MPP(\theta, \mathcal{P}_{\text{dn}}).$$

Starting from distributionally neutral compensation, welfare can be increased by increasing redistribution. Starting from the carbon dividend, welfare can be increased by reducing redistribution. These are local results. They show in which direction welfare can be increased at each benchmark, but do not by themselves establish where the global optimum lies relative to the two benchmarks. In the following sections, we study this question quantitatively by solving directly for the optimal policy, while also allowing for income effects and non-homothetic demand. The quantitative results below confirm the pattern suggested by these directional results: the optimal rebate increases with income, but less steeply than carbon tax payments.

6 Measurement and Quantitative Model

We quantify the model for Germany in 2018. The quantitative analysis relaxes the simplifying assumptions used for the analytical comparative statics, in particular allowing for income effects in labor supply and non-homothetic demand. In Section 6.1, we introduce the data that we use. In Section 6.2, we first construct the carbon-related empirical objects that discipline incidence and the carbon redistribution wedge directly from detailed expenditure and emissions data. We then embed these objects in a structural model in Section 6.3.

6.1 Data

We draw on three primary data sources: (i) German income and consumption expenditure survey (*Einkommens- und Verbrauchsstichprobe*; EVS), (ii) environmentally-extended multi-regional input-output tables from EXIOBASE, and (iii) German administrative tax records (*Lohn- und Einkommensteuerstatistik*; LESt). We briefly describe each dataset below; further details are provided in Appendix B.2 for EXIOBASE and Appendix C for the micro data.

EVS. The *Einkommens- und Verbrauchsstichprobe* (EVS) is a representative cross-sectional household survey that provides detailed data on consumption, income and socioeconomic characteristics. We use the 2018 wave to measure household spending across different goods and services along the income distribution. The survey reports consumption expenditures across 107 categories, classified according to the COICOP system. For the category-level decomposition of carbon footprints and the structural model, we additionally aggregate the 107 goods and services into nine broad consumption categories: clothing, electricity, food, heating, household and personal care, housing, leisure, services, and transport and vacation. In addition, we use detailed information on transfer payments to calibrate a transfer function with a gradual phase-out. Finally, we convert the data from the household level to the tax unit level, resulting in a sample of around 26 000 tax units. We assign each EVS tax unit to the corresponding percentile of the gross labor income distribution from the LESt and compute percentile level moments as survey-weighted averages within these bins.

EXIOBASE. We use the environmentally-extended multi-regional input-output tables from EXIOBASE to compute indirect consumption-based emission intensities. EXIOBASE links economic activity to environmental accounts, including greenhouse gas emissions, through a detailed global input-output framework. We use the product dimension of EXIOBASE, which distinguishes $P = 200$ products and encompasses $K = 49$ regions. To align these data with the COICOP classified consumption data in the EVS, we construct a crosswalk from EXIOBASE to COICOP products. This provides the indirect component of the product-level emission intensities for the 107 COICOP goods. Direct household emission intensities are constructed separately below.

LESt. The *Lohn- und Einkommensteuerstatistik* (LESt) is a 10% cross-sectional random sample of German taxpayers based on administrative income tax records. The unit of observation is the tax unit, defined as either an individual including single persons and married couples filing separately, or a married couple filing jointly. We use data from the 2018 wave on annual gross labor income to calibrate a gross income distribution. In addition, we use total income tax payments to calibrate the status-quo tax system. Our sample includes around 2.28 million tax units.

6.2 Carbon Emission Intensities and Household Carbon Footprints

We first construct product-level emission intensities, household carbon footprints, and the marginal propensity to pollute directly from the detailed COICOP classification. At this stage we do not impose the coarser aggregation into nine broad consumption categories used in the structural model.

6.2.1 Product-Level Emission Intensities

We compute carbon emission intensities for $N = 107$ different COICOP goods and services using a consumption-based approach that attributes emissions to final household demand. For each product, the total emission intensity is the sum of a direct and indirect component. Direct emissions arise at the stage of household consumption (e.g. driving a gasoline-powered car), whereas indirect emissions occur throughout the production chain (e.g. production of a car).

Indirect emission intensities. We begin by computing indirect emission intensities at basic prices using EXIOBASE. Let $\mathbf{L} \in \mathbb{R}^{K \cdot P \times K \cdot P}$ denote the multi-regional Leontief inverse and let $\mathbf{E} \in \mathbb{R}^{1 \times K \cdot P}$ denote the row vector of product-region emission coefficients, where each entry gives the amount of carbon dioxide equivalent (CO_2e) emitted per monetary unit of output for a given product-region pair. To capture the regional composition of German household demand, we construct a sparse allocation matrix $\tilde{\mathbf{Y}} \in \mathbb{R}^{K \cdot P \times P}$. For each EXIOBASE product p, the p-th column of $\tilde{\mathbf{Y}}$ allocates one euro of German household final demand for product p, valued at basic prices, across regions according to the expenditure shares observed in EXIOBASE.[19] Indirect emission intensities at the EXIOBASE product level are then given by

$$\mathbf{f}_{\mathbf{ind}}^{\mathbf{EXIO}} = \mathbf{E} \cdot \mathbf{L} \cdot \tilde{\mathbf{Y}} \quad \in \mathbb{R}^{1 \times P},$$

where $f_{\text{ind},p}^{\text{EXIO}}$ measures the total upstream CO_2e emissions embodied in one euro of German household demand for EXIOBASE product p. We then convert these intensities from basic to purchaser prices and map them from the EXIOBASE product classification into the COICOP classification to make the EXIOBASE based intensities consistent with the household expenditure data. This yields $\tilde{\mathbf{f}}_{\mathbf{ind}}^{\mathbf{COICOP}} \in \mathbb{R}^{1 \times N}$, the vector of indirect emission intensities at purchaser prices in the COICOP classification. Appendix B.3.2 and Appendix B.3.3 provide the details of the conversion from basic to purchaser prices and the mapping from the EXIOBASE to COICOP classification, respectively.

Direct emission intensities. We next compute direct emission intensities at the COICOP level using fuel specific physical emission factors. For each underlying fuel, we convert the physical emission factor into a monetary emission intensity using the corresponding 2018 consumer price and then aggregate these fuel specific intensities to the COICOP level using consumption weighted fuel shares from the 2013 wave of the German Residential Energy Consumption Survey (GRECS). This yields the vector of direct emission intensities, $\tilde{\mathbf{f}}_{\text{dir}}^{\text{COICOP}} \in \mathbb{R}^{N \times 1}$, where $\tilde{f}_{\text{dir},n}^{\text{COICOP}}$ measures the direct CO_2e emissions associated with one euro of purchaser price expenditure in COICOP product n. Appendix B.3.4 provides details on the transformation of physical emission factors into monetary emission intensities.

Total emission intensity. The total emission intensity of COICOP product n is the sum of its indirect and direct components:

$$f_n = \tilde{f}_{\text{ind},n}^{\text{COICOP}} + \tilde{f}_{\text{dir},n}^{\text{COICOP}}.$$

[19] Note that $\tilde{\mathbf{Y}}$ is a vertical stack of K diagonal $P \times P$ blocks, where the k-th block places region k's demand shares across all P products on its diagonal. Appendix B.3.1 describes the construction of $\mathbf{E}$, $\mathbf{L}$ and $\tilde{\mathbf{Y}}$.

The resulting product level emission intensities are reported in Table 7 in Appendix C. As a validation check, combining these intensities with aggregate consumption expenditures implies aggregate emissions of 731.26 $MtCO_2e$ for Germany in 2018, close to the corresponding benchmark reported by the German Statistical Office (German Statistical Office 2025b). See Appendix B.4 for details.

6.2.2 Carbon Footprints and the Marginal Propensity to Pollute

We construct household carbon footprints by combining the COICOP level emission intensities with consumption expenditures from the EVS. Let $e_n(p)$ denote average annual consumption expenditure of percentile p on COICOP product n. The carbon footprint of percentile p is

$$f(p) = \sum_{n=1}^{N} e_n(p) f_n,$$

which captures the direct and indirect emissions embodied in the percentile's consumption basket.

Because our model is static, it abstracts from the intertemporal allocation of consumption across the life-cycle. In the model, disposable income is fully allocated to current consumption, whereas observed expenditures in the data need not equal disposable income because households may save, borrow, or dissave. We therefore apply a 'savings-adjustment': for each percentile, we assign the difference between disposable income and observed expenditure the average emission intensity of that percentile's observed consumption basket. Equivalently, the adjustment rescales total expenditure to disposable income while holding the composition of the observed consumption basket fixed. This reduced-form approach allows us to abstract from the timing of consumption over the life cycle.[20]

Figure 3 shows the carbon footprint, marginal propensity to pollute and emission intensity along the gross labor income distribution. Markers denote the percentile-level moments, while the solid and dashed lines are cubic spline fits. Panel (a) illustrates the carbon footprint in the data in blue and the 'savings-adjusted' carbon footprint in red. Carbon footprints increase with income, but at a decreasing rate, consistent with the literature on environmental Engel curves and household carbon footprints (Levinson and O'Brien 2019, Sager 2019, Hardadi, Buchholz, and Pauliuk 2021, among others). The savings adjustment lowers carbon footprints at the bottom of the income distribution, where observed expenditures exceed measured disposable income, and raises them for higher income tax units with positive saving. Quantitatively, the 90th percentile emits 7.62 times more than the 10th percentile after the savings adjustment. Panel (b) illustrates the corresponding emission intensity of consumption expenditure in red and marginal propensity to pollute out of expenditure in blue, both measured in $kgCO_2e$ per 1 000

[20]This adjustment is similar to the one applied in Craig, Lloyd, and Moore (2025).

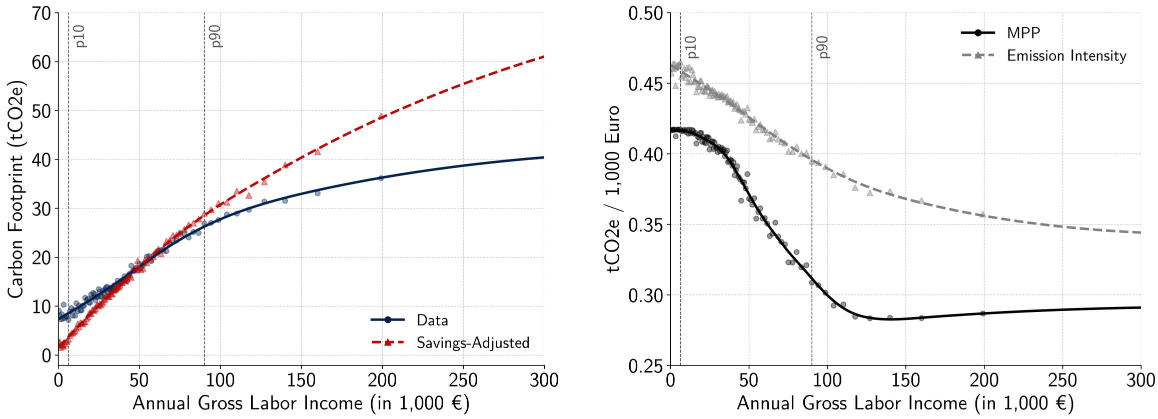

(a) Carbon Footprint

(b) MPP and Emission Intensity

Figure 3: Carbon Footprint, Marginal Propensity to Pollute and Emission Intensity

Notes: The left panel illustrates the carbon footprint $f(p)$ (blue dots) and the savings-adjusted carbon footprint (red triangles) for each income percentile p, both expressed in tCO_2e. The savings adjustment holds the composition of the consumption basket fixed and rescales all expenditures by a common factor such that adjusted total expenditure equals disposable income. The right panel shows the marginal propensity to pollute (MPP, blue dots) and the emission intensity (red triangles), both expressed in tCO_2e per 1 000 Euro. The MPP is defined as the derivative of the spline-fitted carbon footprint with respect to total expenditure, evaluated at each percentile: $\text{MPP}(p) = \hat{f}'(e(p))$, where $\hat{f}(\cdot)$ is a cubic smoothing spline of $f(p)$ on $e(p)$. The emission intensity is defined as $\iota(p) = \hat{f}(p)/e(p)$. In both panels, the lines show cubic smoothing spline fits to the respective data moments. The x-axis shows annual gross labor income of a tax unit (in 1,000 EUR). The grey vertical dashed lines correspond to the 10th and 90th percentile of the annual gross labor income distribution.

Euros of expenditure. The MPP is computed as the local slope of the carbon Engel curve, $MPP(p) = \hat{f}'(e(p))$, where $\hat{f}(e(p))$ is a cubic smoothing spline fitted to the percentile-level relationship between the carbon footprint $f(p)$ and total expenditure $e(p)$.[21] The emission intensity is defined as $\iota(p) = \hat{f}(p)/e(p)$. Quantitatively, the MPP ranges from approximately 0.42 to 0.29 $kgCO_2e$ per EUR of expenditure and the emission intensity from 0.46 to 0.35 $kgCO_2e$ per EUR of expenditure.[22] Both measures decline along the gross labor income distribution. This reflects the non-homotheticity of consumption: as income rises, spending shifts toward less carbon-intensive goods and services, so an additional euro of expenditure generates fewer emissions than the average euro spent at lower income levels.[23]

[21] The local slope of the carbon Engel curve, estimated from cross-sectional data, may reflect both income effects and systematic heterogeneity in preferences correlated with expenditure. Empirical evidence from Craig, Lloyd, and Moore (2025) supports interpreting this slope as a reasonable proxy for the causal marginal response. In a related setting, they find that cross-sectional Engel slopes for carbon-intensive goods are very close to elicited within-household responses to a small permanent income increase, implying little residual heterogeneity.

[22] Berr, Kindermann, and Le Blanc (2026) provide a complementary life-cycle benchmark of the MPP. Using a structural model, they obtain a model-implied lifetime income emissions slope of 0.54 $kgCO_2e$ per euro. Their estimate is an average lifetime gradient across 26 EU countries, whereas ours is a local expenditure derivative for Germany. The closest counterpart to their income-based measure is therefore our savings-adjusted MPP shown in Figure 5(b), which is higher than the expenditure-based MPP.

[23] This is also consistent with evidence from Aghion, Boppart, Peters, Schwartzman, and Zilibotti (2025) for the United States that demand shifts toward relatively clean, service-intensive goods as income rises.

We further decompose the MPP into a spending-level and a composition component:

$$MPP(e) = \iota(e) + e\,\frac{d\iota(e)}{de}.$$

The spending-level component captures the emissions generated by an additional euro of expenditure holding the consumption basket fixed and coincides with the average emission intensity $\iota(e)$. The composition component, $e\,d\iota/de$, captures change in emissions arising from changes in the composition of the consumption basket as total expenditure rises. In Figure 3(b), this component is simply the vertical difference between the MPP and the average emission intensity.[24] Because spending shifts toward less carbon-intensive goods as expenditure rises ($d\iota/de < 0$), the composition component is negative, and the total MPP lies below the spending-level component throughout most of the income distribution: shifts in marginal spending attenuate the emissions generated by an additional euro. This attenuation is strongest at low and middle incomes; above roughly EUR 120 000 the reallocation of marginal spending weakens and the total MPP converges toward the spending-level component. Figure 9 in Appendix D.1 illustrates the composition component.

6.2.3 Anatomy of the Income Gradient in Carbon Footprints

As just described, Figure 3 summarizes the main empirical objects that discipline the distributional effects of carbon pricing in the model. The level of the carbon Engel curve determines carbon tax liabilities across the income distribution, while its local slope with respect to expenditure disciplines the carbon redistribution wedge. We next examine the sources of this income gradient. Since a tax unit's carbon footprint aggregates expenditures across goods and emissions sources, the carbon-income gradient can reflect several distinct margins: the distinction between emissions released directly by households and emissions embodied in production, changes in the observable composition of tax units along the income distribution, and changes in consumption across product categories. We examine each margin in turn.

Direct vs. indirect emissions. We first split total carbon footprints into direct emissions and indirect emissions. Indirect emissions account for most of the income gradient, whereas direct emissions are substantially flatter. This decomposition clarifies the scope of the benchmark policy. As described in Section 3, our model implements a consumption-based carbon price whose tax base is the household's carbon footprint, including both direct emissions and indirect emissions embodied in purchased goods and services. By contrast, a policy that prices only direct household emissions would use the direct footprint $f^{dir}(e)$ and the corresponding

[24] Total differentiation of the carbon footprint $f(e) = \iota(e)\,e$ with respect to expenditure e gives $MPP = df/de = \iota(e) + e\,d\iota(e)/de$. In the data, we measure the spending-level component as the observed average emission intensity $\iota(e(p)) = f(p)/e(p)$ at each percentile p, and recover the composition component $e(p)\,d\iota(e(p))/de$ as the residual $\widehat{MPP}(p) - \iota(e(p))$.

$MPP^{dir}(e)$ in place of the total footprint $f(e)$ and total $MPP(e)$. Since the direct component is flatter, such a narrower policy would imply a smaller carbon tax liability gradient and, for a given carbon price, a smaller carbon redistribution wedge. Figure 10 in Appendix D.2 illustrates the carbon footprints decomposed into direct and indirect emissions along the gross labor income distribution.

Horizontal heterogeneity. We next consider observable horizontal heterogeneity. Carbon footprints may vary not only with income, but also across tax units with similar income, reflecting differences in household structure, demographics, housing characteristics, and location. If these characteristics vary systematically with income, the income gradient combines vertical carbon inequality with observable composition effects. To quantify this component, we construct covariate-standardized carbon footprints for each income percentile. Using EVS tax unit microdata, we estimate a linear projection of carbon footprints on a income percentile fixed effects and controls for household size, education, homeownership, heating fuel, rural residence, and age. We then recover the covariate-standardized percentile profile from the estimated percentile fixed effects, evaluating the controls at their weighted empirical distribution in the full estimation sample for every percentile. The resulting profile therefore captures how average carbon footprints would vary across the income distribution when the distribution of observed covariates is held fixed across percentiles.[25]

The adjustment attenuates, but does not eliminate, the income gradient in carbon footprints. Holding the distribution of observed covariates fixed across percentiles, the MPP is about 23% lower on average, but remains positive throughout the income distribution and declines over most of the income distribution. Since the carbon redistribution wedge is proportional to the MPP, taking out observed horizontal heterogeneity reduces the quantitative scope for additional redistribution, but leaves the qualitative mechanism unchanged. Figure 11 in Appendix D.2 illustrates the covariate-standardized carbon footprints and the corresponding MPP along the gross labor income distribution.

Product category decomposition. Finally, using the nine broad categories introduced in Section 6.1, we decompose the percentile-level footprint by product category. For each income decile, we compute the average product-specific carbon footprint. The income gradient is present across the consumption basket: category-specific footprints increase with income in each category. Transport and vacation is the largest single contributor, accounting for 41.2% of the gap in carbon emissions between the 10th and 90th percentiles. The remaining categories jointly account for 58.8%. Hence, the income gradient in carbon emissions is not an

[25] Note that this exercise is purely descriptive. It removes the component of the income gradient associated with observed covariates, but it does not account for unobserved heterogeneity or isolate a causal effect of income on emissions. See Appendix D.2 for details.

artifact of a single high emission product. Figure 12 in Appendix D.2 shows the category level decomposition.

6.3 Structural Model

The preceding subsection constructs the carbon Engel curve and its local slope directly from the detailed COICOP data. Counterfactual carbon pricing additionally requires predicting how households reallocate expenditure when relative prices change. We therefore embed these empirical objects in a structural model with $I = 9$ broad consumption categories, nested non-homothetic CES demand, endogenous abatement, the German tax-transfer system, and the estimated income distribution.

6.3.1 Preferences

For the functional form of $U(\mathbf{c})$, we adopt a utility function with nested non-homothetic CES preferences. At the upper tier, there is a consumption basket $\mathbf{c} = (c_1, c_2, ..., c_I)$. At the lower tier, each good $i = 1, \ldots I$ is a non-homothetic CES composite of a 'brown' and 'green' variety, capturing substitution possibilities between goods with different carbon emission intensities. The nested non-homothetic CES structure allows for varying expenditure shares along the income distribution in both goods (upper tier) and green and brown varieties (lower tier).

Upper-tier utility function. We consider non-homothetic CES preferences defined over a consumption basket $\mathbf{c} = (c_1, c_2, ..., c_I)$. Following Comin, Lashkari, and Mestieri (2021), the functional form of $U(\mathbf{c})$ is given by[26]

$$U(\mathbf{c}) = \frac{\mathbb{C}(\mathbf{c})^{1-\gamma}}{1-\gamma}, \tag{31}$$

where $\gamma \geq 0$ denotes the coefficient of relative risk aversion and $\mathbb{C}$ is a consumption aggregator implicitly defined by

$$\sum_{i=1}^{I} (\Omega_i \mathbb{C}(\mathbf{c})^{\varepsilon_i})^{\frac{1}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} = 1 \tag{32}$$

with $\sigma > 0$, $\Omega_i > 0\ \forall i$. We consider the case $\sigma < 1$, where these broader consumption categories are complements. In this case, $\varepsilon_i > 0\ \forall i$. With non-homothetic preferences, the elasticity of substitution σ is constant across consumption goods, but the expenditure shares households allocate to different goods vary with their net income $e(\theta)$.

The elasticity of consumption of good i w.r.t. to overall expenditure e is given by

$$\xi_i \equiv \frac{\partial \log c_i}{\partial \log e} = \sigma + (1-\sigma)\frac{\varepsilon_i}{\bar{\varepsilon}}, \tag{33}$$

[26] To be precise, we follow the notation of an older working paper version of this paper (Comin, Lashkari, and Mestieri 2017).

where $\overline{\varepsilon} = \sum_i = \psi_i \varepsilon_i$ and ψ_i is the expenditure share of good i. Note that goods with $\varepsilon_i > \overline{\varepsilon}$ are luxuries, while those with $\varepsilon_i < \overline{\varepsilon}$ are necessities.

Lower-tier utility function. Each good c_i is a composite of a 'brown' and 'green' variety, which differ in their carbon emission intensities. We denote these varieties by c_{ij}, with $j \in \{b, g\}$ indicating the 'brown' (b) and 'green' (g) variety, respectively.

$$\sum_{j \in \{b,g\}} (\Omega_{ij} c_i^{\varepsilon_{ij}})^{\frac{1}{\sigma_i}} c_{ij}^{\frac{\sigma_i - 1}{\sigma_i}} = 1,$$

where σ_i denotes the elasticity of substitution between two varieties and $\Omega_{ij} > 0$. For the inner nest, we consider both cases: $\sigma_i > 1 \; \forall i$ and $\sigma_i < 1 \; \forall i$ implying $\varepsilon_{ij} < 0 \forall i, j$ and $\varepsilon_{ij} > 0 \forall i, j$ respectively. This preference structure allows for substitution from environmentally harmful to environmentally friendly goods within a consumption category i when the relative price of brown varieties increases due to the introduction of a carbon tax.

6.3.2 Green and Brown Varieties

To map the detailed COICOP data into the lower-tier nests of the model, we classify each product within its broad category as either 'brown' or 'green' based on its emission intensity. Specifically, we apply K-means clustering with $K = 2$ separately within each category, using the log emission intensity as the clustering variable. We label the cluster with the higher average emission intensity as 'brown' and the other as 'green'.[27] We then compute the expenditure-weighted average emission intensity of all goods classified as brown and green within each product category. This yields the emission intensity f_{ij}, where $j \in \{b, g\}$ indicates the brown or green variety of product category i. Table 1 summarizes the classification and the resulting average emission intensities f_{ij}. Note that electricity is not split into a brown or green variety, as this product category comprises only one product.

Before estimating the preference parameters, we specify the remaining quantitative inputs needed for the calibration: the carbon abatement technology, which determines residual emission intensities and consumer prices; the tax-transfer system, which maps gross income into disposable resources; and the income distribution.

[27] Table 7 in Appendix C contains the classification of all 107 products into 'green' and 'brown' varieties. Note that this mapping of the 107 products into the broader categories implies an aggregation error. To illustrate this error, we construct the adjusted empirical measure of the carbon footprint by aggregating each tax unit's EVS expenditures to the corresponding product–variety cells, multiplying these expenditures by the associated expenditure weighted average emission intensities f_{ij} and summing over all cells. Appendix D.3 shows that the approximation error from this aggregation is quantitatively small.

Product Category	**Cutoff Emission Intensity**	**Average Emission Intensity**	
		Brown	Green
Clothing	0.1849	0.2807	0.1199
Electricity	–	1.7913	1.7913
Food	0.4905	0.631	0.3944
Heating	5.9558	8.6185	3.8923
Household and Personal Care	0.2003	0.3001	0.1498
Housing	0.1416	0.3641	0.059
Leisure	0.1741	0.2571	0.1229
Services	0.1703	0.2737	0.1063
Transport and Vacation	0.4761	1.5732	0.218

Table 1: Emission intensities of Brown and Green Varieties

Notes: Emission intensities are measured in kgCO$_2$e per euro of expenditure. Within each product category, products are classified into brown and green varieties using K-means clustering with $K = 2$ applied to product-level emission intensity. The second column reports the category-specific cutoff implied by this procedure. Columns 3 and 4 report the corresponding expenditure-weighted average emission intensities. Electricity is not split into brown and green varieties because this category contains only one product.

6.3.3 Carbon Abatement Technology

We calibrate the abatement technology following the DICE literature.[28] For product i and variety j, abatement costs are given by

$$\kappa_{ij}(s_{ij}) = \bar{f}_{ij} \frac{k}{1+\psi} s_{ij}^{1+\psi} \quad \forall i, j,$$

where $\bar{f}_{ij}$ denotes the emission intensity absent abatement and $s_{ij} \equiv a_{ij}/\bar{f}_{ij} \in [0,1]$ is the abatement share. The scale parameter k is the backstop price, i.e. the carbon tax that induces full abatement, while ψ governs the curvature of abatement costs and equals the inverse of the elasticity of the abatement share with respect to the carbon tax. Both parameters are common across goods and varieties. The implied optimal abatement share at carbon price τ_{co2} is

$$s_{ij}(\tau_{\text{co2}}) = \min\left\{ \left(\frac{\tau_{\text{co2}}}{k}\right)^{1/\psi}, 1 \right\} \quad \forall i, j. \tag{34}$$

Residual emission intensities are therefore

$$f_{ij}(\tau_{\text{co2}}) = \bar{f}_{ij} \left[1 - s_{ij}(\tau_{\text{co2}})\right] \quad \forall i, j. \tag{35}$$

[28]We use the abatement cost schedule to discipline the responsiveness of residual emission intensities to carbon prices. We map this aggregate mitigation margin into the emissions-intensity adjustment in our model, while changes in household consumption imply emission reductions on top.

Following DICE-2023, we set $\psi = 1.6$ and $k = 600$ EUR/tCO_2e.[29] We then infer $\bar{f}_{ij}$ from the observed 2018 emission intensity in the data. Given the average effective carbon price on emissions embodied in German household consumption in 2018, $\tau_0 = 3.91$ EUR/tCO_2e[30], (34) and (35) then imply

$$s_0 = \left(\frac{\tau_0}{k}\right)^{1/\psi} \qquad \text{and} \qquad \bar{f}_{ij} = \frac{f_{ij}}{1 - s_0}.$$

6.3.4 Tax-Transfer System

We adopt a parametric specification for the tax-transfer system proposed by Ferriere, Grübener, Navarro, and Vardishvili (2023). Total tax payments are given by

$$\mathcal{T}(y) = \exp\left[\log(\lambda_{tax})(y/\bar{y})^{-2\tau_{tax}}\right] y - \mathcal{T}_0(y),$$

where $\bar{y}$ denotes average gross labor income. The first term specifies a two-parameter tax schedule: the parameter λ_{tax} governs the overall level of taxation and the parameter τ_{tax} controls the degree of progressivity. The transfer function $\mathcal{T}_0(y)$ phases out with labor income and is given by

$$\mathcal{T}_0(y) = m\bar{y}\frac{2\exp\left[-\zeta\frac{y}{\bar{y}}\right]}{1 + \exp\left[-\zeta\frac{y}{\bar{y}}\right]},$$

where m captures the size of transfers given to households with zero income as a share/multiple of average labor income $\bar{y}$ and ζ determines the phase-out rate of transfers. This specification allows us to jointly fit the tax payments of the bottom and the top of the income distribution.

We estimate the tax parameters $(\lambda_{tax}, \tau_{tax})$ using tax return data from the LESt, and the transfer parameters (m, ζ) using data on transfers from the EVS. We estimate the parameters using non-linear least squares. Table 2 presents the estimated parameters and Figure 4(a) illustrates the marginal rates obtained from the estimated tax-transfer function. The solid red line depicts the statutory marginal income tax rate, while the dashed gray line shows the transfer phase-out rate. Transfers decline steeply with gross labor income and are fully phased out at an annual gross labor income of around €40 000. The effective marginal income tax rate is the sum of the statutory marginal income tax rate and the transfer phase-out rate. As a result, the EMTR exhibits a sharp drop at the point where transfers are fully phased out, after which it converges to the statutory marginal tax rate. Appendix E.2 provides further details on the estimation process.

[29] Barrage and Nordhaus (2024, Appendix E) report a DICE-2023 backstop price of 695 USD/tCO_2 in 2020, measured in 2019 dollars. Backcasting this value to 2018 using DICE's assumed 1% annual decline gives 709 USD/tCO_2. Using the 2018 average market exchange rate of 1.181 USD per euro gives about 600 EUR/tCO_2. In Section 7.3, we consider $\psi = 1$ and $\psi = 2$ as robustness exercises.

[30] We approximate the average effective carbon price on emissions embodied in German household consumption as the uniform carbon price that would generate the same aggregate carbon cost as the EU ETS component embodied in the 2018 consumption basket. See E.1 for details.

6.3.5 Income Distribution

To calibrate the distribution of gross labor income, we use administrative income tax records from the LESt for the year 2018. Gross labor income is defined as the sum of income from employment and self-employment. We estimate the income distribution $\tilde{h}(y)$ non-parametrically using a standard kernel density estimation. Figure 4(b) illustrates the Local Pareto parameter $\tilde{h}(y)y/(1-\tilde{H}(y))$ derived from the estimated income distribution. At the top, the income distribution $\tilde{h}(y)y/(1-\tilde{H}(y))$ converges to roughly 2, consistent with a Pareto tail with a Pareto parameter of around 2. Appendix E.3 documents the estimation procedure in detail.

Given the estimated income distribution and the current German tax-transfer system, we infer the skill distribution $h(\theta)$ by inverting the labor supply first-order condition, following Saez (2001).

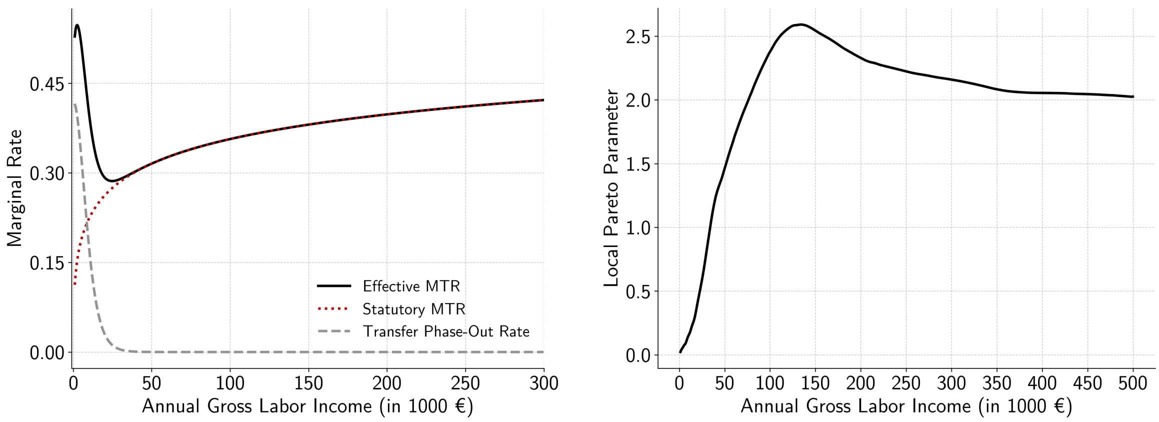

(a) Tax-Transfer System

(b) Income Distribution

Figure 4: Tax-Transfer System and Income Distribution

Notes: The left panel illustrates the statutory marginal income tax rate (dotted red line) and the transfer phase-out rate (dashed gray line) based on our estimates of the parametric function for the tax-transfer system proposed by Ferriere, Grübener, Navarro, and Vardishvili (2023). The effective marginal income tax rate (solid black line) is the sum of the statutory marginal income tax rate and the transfer phase-out rate. The right panel displays the local Pareto parameter $h_y(y)y/(1-H_y(y))$, where $H_y(y)$ is the cumulative distribution function and $h_y(y)$ is the probability density function of the gross labor income distribution. The density is estimated using a standard kernel density estimation. The x-axis shows gross labor income y of a tax unit.

6.3.6 Calibration of Preferences

Given the product aggregation and the auxiliary quantitative inputs described above, we now calibrate the parameters governing household preferences. The calibration of the utility function, as represented by (31)–(33), requires a total of 43 parameters: $\{\Omega_{ig}\}_{i=1}^{8}$, $\{\epsilon_{ig}\}_{i=1}^{8}$ and $\{\sigma_i\}_{i=1}^{8}$ for the lower tier; $\{\Omega_i\}_{i=1}^{8}$, $\{\epsilon_i\}_{i=1}^{8}$ and σ for the upper tier. Finally, the more standard Frisch elasticity parameter φ and the curvature parameter γ, which would equal the coefficient of relative risk aversion case in case of homothetic preferences.[31] Next, we explain which of those parameters we calibrate externally and which internally.

Externally calibrated parameters. We fix the Frisch elasticity of labor supply at $\varphi = 0.5$, in line with empirical estimates (Chetty, Guren, Manoli, and Weber 2011). We set the curvature parameter of consumption utility to $\gamma = 0.5$ and the elasticity of substitution for the upper tier utility function to $\sigma = 0.3$ (Comin, Lashkari, and Mestieri 2021). Finally, we assume that the elasticity of substitution of the lower-tier utility function, i.e. between 'brown' and 'green' varieties, to be identical across all products and set to $\sigma_i = 1.5$. Table 2 summarizes the set of externally calibrated parameters.[32]

Parameter	Description	Value
	Utility Function	
φ	Frisch elasticity of labor supply	0.5
γ	Coefficient of relative risk aversion	0.5
σ	Elasticity of substitution between products	0.3
$\sigma_i\ \forall i$	Elasticity of substitution between varieties	1.5
	Tax-Transfer System	
λ	Overall level of taxation	0.25
τ	Tax progressivity	0.08
m	Transfers given to households with zero income	0.1
ζ	Transfer phase-out rate	8.4
$\bar{y}$	Average gross labor income (in €)	42 189

Table 2: Externally Calibrated Parameters

Notes: The table reports the set of externally calibrated parameters used in the model. The tax-transfer system parameters are chosen to match the German tax-transfer system in 2018, based on data from the LESt and EVS.

Internally calibrated parameters. We calibrate the remaining preference parameters in two steps, exploiting how expenditure patterns in the EVS vary along the gross labor income distribution.

First, for each product i, we calibrate the lower-tier parameters governing the allocation between the 'brown' and 'green' varieties. We use the 'brown' variety as the reference and impose the normalization $\Omega_{ib} = 1$ and $\epsilon_{ib} = -1$ for all i. The parameters $\{\Omega_{ig}, \epsilon_{ig}\}_{i=1}^{I-1}$ on the 'green' variety are then chosen by minimum distance to match the empirical within-product expenditure shares on the 'green' variety across income deciles. Intuitively, the parameter Ω_{ig}

[31] Note that electricity has no lower-tier utility nest because this category contains only a single product and is therefore not differentiated into brown and green varieties.

[32] The chosen values for γ, σ and φ allow the model to replicate reasonable values of the income effect $\eta(\theta)$. Figure 15 in Appendix E.4.1 illustrates the income effect implied by the model. In Section 7.3 we provide robustness for these parameter choices.

disciplines the average green share within a category, while the lower-tier non-homotheticity parameter ϵ_{ig} disciplines how this share varies with income.

Second, conditional on the calibrated lower-tier nests, we calibrate the upper-tier parameters $\{\varepsilon_i, \Omega_i\}_{i=1}^{I-1}$, which govern the allocation of total expenditure across product categories. We normalize the parameters of one reference category and choose the remaining parameters Ω_i and non-homotheticity parameters ϵ_i by minimum distance, targeting the product-level expenditure shares across income deciles. The average level of a product's budget share primarily identifies Ω_i, while the income gradient of the budget share identifies the non-homotheticity parameter ε_i. The internally calibrated parameters are reported in Table 8 in E.4.2. See Appendix E.4.2 for details on the internal calibration strategy.

6.3.7 Model Fit

We first validate the fit of the calibrated demand system to the expenditure share moments used in the internal calibration. As shown in Figure 16 in Appendix E.4.2 the model closely tracks the empirical income gradient in both the allocation of expenditure across products and the within-category allocation between 'brown' and 'green' varieties.

Second, we evaluate the model's fit to how carbon emissions vary with income. Figure 5(a) illustrates the model-implied carbon footprints along gross labor income (solid red line), together with the corresponding empirical moments from the EVS (grey lines). For the empirical moment, we apply the savings adjustment and aggregate goods into the same broad category–variety classification as in the model. The model captures well the overall income gradient in carbon emissions, but implies a less concave relationship than the data.[33] The model-implied carbon footprint of households at the 90th percentile is 7.92 times higher than of those at the 10th percentile, compared with 7.91 in the savings-adjusted aggregated data. Figure 5(b) shows the model-implied MPP and emission intensity profiles (solid lines), together with the corresponding empirical moments from the EVS (dashed lines). The model captures the declining pattern across the income distribution well, although the model-implied values are slightly higher than the empirical moments.

6.4 Welfare Criterion and Inverse-Optimum Weights

To isolate how the carbon constraint changes optimal redistribution, we hold fixed the welfare criterion that rationalizes the estimated 2018 tax–transfer system. We calibrate this criterion using an inverse-optimum approach for the calibrated model just described in Section 6.3. Evaluating the optimal tax condition in Appendix A.2 (which is the more general version

[33] This partially reflects that utility parameters are disciplined by decile-level expenditure-share moments, which target average expenditure patterns within income bins.

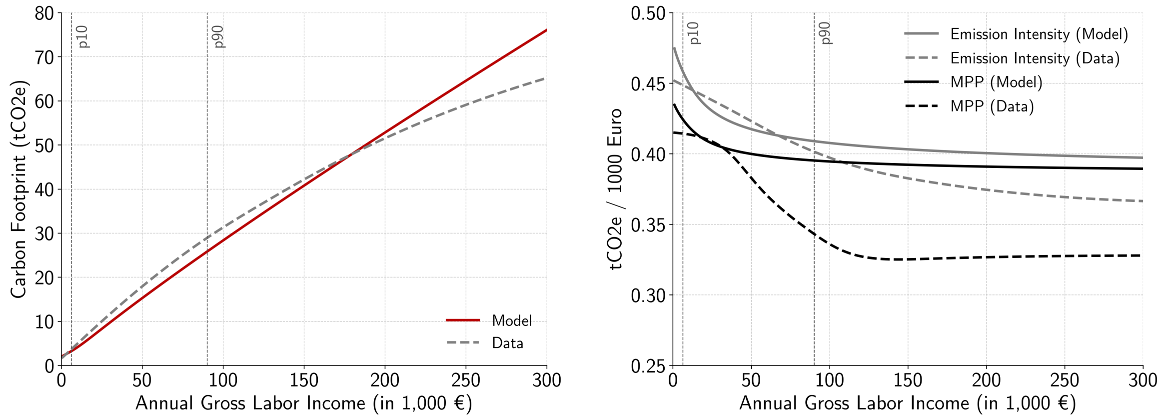

(a) Carbon Footprint (b) Emission Intensity

Figure 5: Model Fit

Notes: The left panel plots the model-implied carbon footprint (red solid line) against a cubic smoothing spline fitted to the savings-adjusted carbon footprint estimated with EVS data (grey dashed line). The right panel shows the model-implied average emission intensity, per Euro of consumption expenditure (grey solid line) and the marginal propensity to pollute (black solid line), both measured in tCO_2e per 1 000 Euro. The MPP is defined as the derivative of the carbon footprint with respect to total expenditure, and emission intensity is defined as the ratio of the carbon footprint to expenditure. In both figures, the corresponding savings-adjusted empirical counterparts from the EVS are shown as dashed lines. To make the model and data comparable, the EVS based spline applies the same broad category–variety aggregation as the model and rescales expenditures to match disposable income. The x-axis shows gross labor income y of a tax unit. The grey vertical dashed lines correspond to the 10th and 90th percentiles of the annual gross labor income distribution.

of Lemma 2 without Assumption 1 being imposed) recovers the status-quo profile of social marginal welfare weights $g(\theta, P_{\text{sq}})$ up to a common normalization.

The inverse-optimum condition recovers $g(\theta, P_{\text{sq}})$, but counterfactual policy analysis additionally requires specifying how social marginal welfare weights evolve when household utility and the marginal utility of expenditure change. Conditional on a choice of the transformation $G(V)$, we therefore infer the Pareto weights from

$$\omega(\theta) \propto \frac{g(\theta, P_{\text{sq}})}{G'(V(\theta, P_{\text{sq}}))\, u_e(\theta, P_{\text{sq}})},$$

where $V(\theta, P_{\text{sq}})$ and $u_e(\theta, P_{\text{sq}})$ follow directly from the calibrated structural model. Finally, we normalize the Pareto weights such that

$$\int_{\underline{\theta}}^{\bar{\theta}} \omega(\theta)\, dH(\theta) = 1.$$

In the benchmark specification, we assume $\mathcal{G}(V) = V$. The implied social marginal utilities and Pareto weights are shown in Figure 17 in Appendix E.5. Within each parameter specification, we hold G and the inferred Pareto weights $\omega(\theta)$ fixed when comparing the status quo with the carbon-cap counterfactual.

7 Quantitative Policy Analysis

We now quantify how a binding cap on aggregate emissions reshapes redistribution in the calibrated 2018 German economy. For a given carbon cap, the government chooses the carbon price and an income-dependent rebate schedule, holding fixed the welfare criterion. Although the benchmark sets $G(V) = V$, the carbon-dividend result in Proposition 4 does not apply because the quantitative model imposes neither Assumption 1 nor Assumption 3: consumption utility is curved and non-homothetic, and the estimated status-quo tax-transfer schedule is non-affine.

First, we consider a cap that implies a 25% reduction in carbon emissions as the baseline in Section 7.1. In Section 7.2 we then analyze the robustness of our results along the abatement path to zero emissions. Section 7.3 considers various robustness checks.

7.1 Baseline Results: 25% Carbon Emissions Reduction

Our benchmark imposes a 25% reduction in aggregate model emissions relative to the 2018 baseline, $\bar{F} = 0.75F_{2018}$. We adopt this target as a conservative policy benchmark. Reductions implied by current German and European climate targets imply reductions between 27% and 49% relative to 2018.[34] However, these policy targets are not directly comparable to the model cap. They apply to territorial emissions, whereas the model prices consumption-based household emissions, including both direct and indirect emissions. Section 7.2 studies the full abatement path to zero emissions.

Carbon tax. The optimal carbon tax amounts to $\tau_{co2} = 59.91$ €/tCO2e.[35] To gauge the fiscal scale, we express carbon tax revenue in VAT-equivalent terms. Holding the VAT base fixed, the carbon tax required to implement the benchmark cap raises the same revenue as a 2.15 percentage point increase in the statutory VAT rate.[36] This highlights the substantial revenue potential and motivates the question of how these revenues should be rebated.

[34] In particular, the 2030 target under the German Federal Climate Change Act (*Bundes-Klimaschutzgesetz*, KSG) corresponds to a reduction of about 49% relative to 2018. The 2030 target under the European Union Effort Sharing Regulation (ESR) implies a reduction of about 45% relative to 2018. The target under the European Union's second emissions trading system for buildings and road transport (ETS2) corresponds to a reduction of about 27% relative to 2018. Appendix E.6 provides the underlying calculations.

[35] Throughout this section, all reported carbon prices and carbon tax revenues are expressed relative to the 2018 baseline: carbon prices are reported in excess of the effective baseline price $\tau_0 = 3.91$ €/tCO2e, and carbon tax revenues are reported as additional revenues relative to the baseline carbon tax revenue of the calibrated 2018 economy.

[36] This VAT comparison is purely a reporting device, as VAT is not modeled explicitly in the quantitative analysis. We compare the actual increase in carbon-tax revenue relative to the calibrated 2018 baseline with the revenue from an increase in the statutory VAT rate. The hypothetical VAT base is recovered from purchaser-price consumption using the German standard VAT rate of 19%, treating the carbon-inclusive price as subject to VAT. Appendix F provides details.

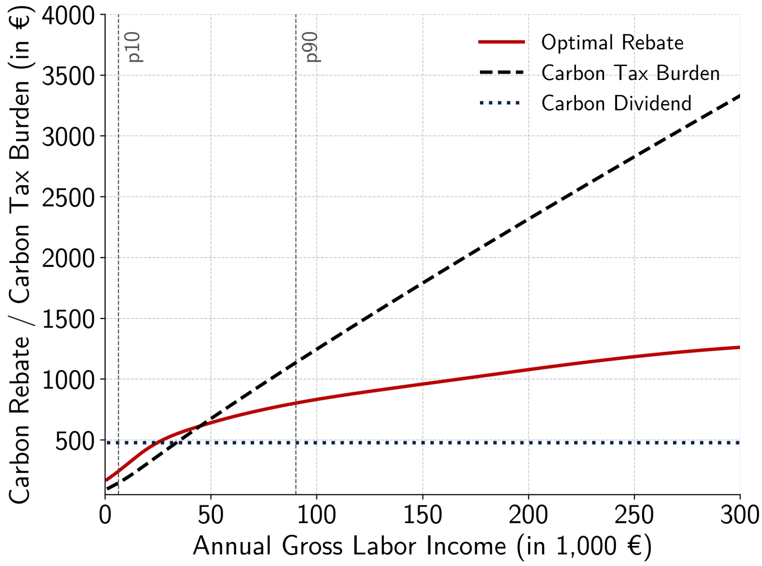

Figure 6: Optimal Carbon Rebate

Notes: The figure shows the optimal carbon rebate as a function of gross labor income (red solid line), the rebate under a carbon dividend policy (blue dotted line) and the carbon tax burden under the optimal carbon rebate (black dashed line). The grey vertical dashed lines correspond to the 10th and 90th percentiles of the annual gross labor income distribution.

Main results. Figure 6 illustrates the optimal income-dependent carbon rebate schedule $R_{\mathrm{CO}_2}(y)$ (red solid line) as a function of gross labor income. The schedule is upward sloping: the annual rebate is about € 200 per tax unit at the bottom of the income distribution and about € 1 200 at an annual gross labor income of € 300 000. To assess the distributional pattern of the rebate policy, we define the regressivity ratio[37] as the ratio of the rebate received by the 90th percentile to that received by the 10th percentile of the income distribution. Formally,

$$RR_{90/10}^{\text{rebate}} \equiv \frac{R_{\mathrm{CO}_2}(y_{90})}{R_{\mathrm{CO}_2}(y_{10})},$$

where y_p denotes annual gross labor income at percentile p, with $y_{90} = 91,909$ and $y_{10} = 6,495$. For the optimal rebate schedule, we find $RR_{90/10}^{\text{rebate}} = 3.31$, indicating that the tax units in the 90th percentile receive a 3.31 times higher rebate than those in the 10th percentile. In this absolute-payment sense, the rebate policy in isolation is regressive. By contrast, the *progressivity ratio* of the carbon tax is $PR_{90/10}^{\text{carbon tax}} = 7.87$, indicating that the 90th percentile pays 7.87 times more in carbon taxes than the 10th percentile. Recall that this latter result directly rests on how carbon footprints vary along the income distribution.

Although the optimal rebate rises with income in absolute amounts, it rises much less steeply than carbon tax payments. As Figure 6 shows, low-income households are overcompensated

[37] We use "regressivity" and "progressivity" here to describe the income gradient in absolute rebate and carbon tax payment amounts, respectively. These ratios are not measures of progressivity in the usual average or marginal rate sense; the progressivity of the overall policy is assessed below using effective average tax rates.

relative to their carbon tax payments, whereas high-income households receive less in rebates than they pay in carbon taxes.

To assess the overall incidence of the reform more thoroughly, Figure 7(a) reports changes in effective average tax rates. The red solid line illustrates the net fiscal burden of the carbon policy, defined as the change in the effective average tax rate, as a function of the percentile of the income distribution.[38] The net burden rises from approximately -10 percentage points at the very bottom to 1 percentage point at the top of the distribution. Tax units up to approximately the 50th percentile are net fiscal beneficiaries of the carbon policy, receiving more in rebates than they pay in carbon taxes.

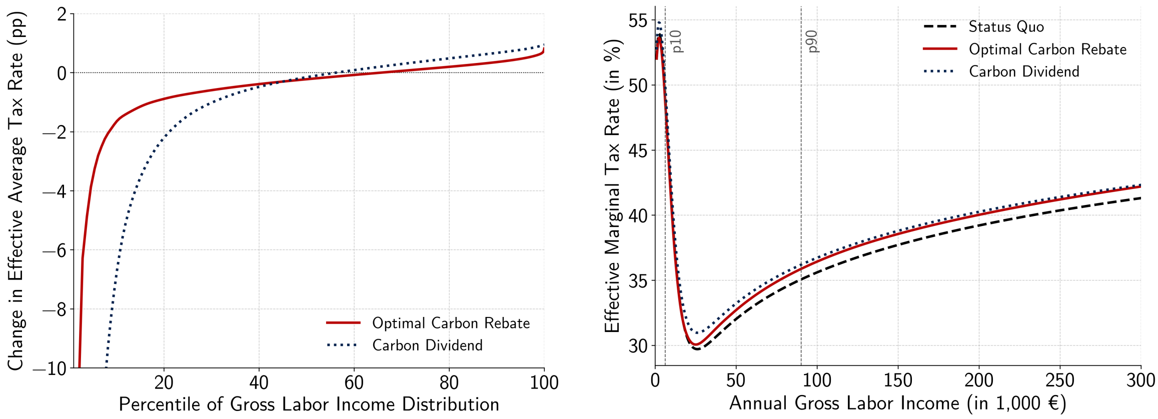

(a) Net Fiscal Burden of Carbon Policy (b) Effective Marginal Tax Rates

Figure 7: Redistributive Implications of Carbon Policy

Notes: The left panel shows the change in the effective average tax rate $\bar{\tau}_{\text{eff}}$ across percentiles of the gross labor income distribution, defined as $\bar{\tau}_{\text{eff}}(\theta) = (\mathcal{T}(y(\theta)) + \tau_{co_2} f(\theta))/y(\theta)$. The right panel shows the EMTR τ_{eff} as a function of gross labor income. The black dashed line represents the status quo without carbon policies, the red solid line the optimal carbon rebate policy, and the blue dotted line the carbon dividend policy. The grey vertical dashed line in the right panel marks the 10th and 90th percentiles of the annual gross labor income distribution.

To summarize the redistributive effect of the carbon policy in a single statistic, we compare effective average tax rates (EATR) at the 10th and 90th percentiles of the gross labor income distribution. For policy $x \in \{\text{sq}, \text{co2}\}$, define the EATR gap as $\Pi^x_{90/10} \equiv \bar{\tau}^x_{\text{eff}}(y_{90}) - \bar{\tau}^x_{\text{eff}}(y_{10})$, where $\bar{\tau}^x_{\text{eff}}(y_p)$ denotes the effective average tax rate evaluated at that income under policy x. We measure the shift in redistribution as the change in the EATR gap:

$$\Delta\Pi_{90/10} \equiv \Pi^{\text{co2}}_{90/10} - \Pi^{\text{sq}}_{90/10}.$$

We find $\Delta\Pi_{90/10} = 2.00$ percentage points, indicating an economically meaningful increase in redistribution under the optimal carbon policy. Table 10 in Appendix F reports the underlying EATRs and implied differences.

[38]In model notation, we define the EATR as $\bar{\tau}_{\text{eff}}(\theta) = \frac{\mathcal{T}(y(\theta)) + \tau_{co_2} f(\theta)}{y(\theta)}$.

Carbon reduction	10%	25%	50%	75%	90%
Carbon tax (€/tCO_2e)	17.52	59.91	172.65	347.19	486.93

Table 3: Carbon tax by emissions-reduction target

Notes: This table illustrates the different levels of carbon taxes needed to achieve a carbon reduction of the indicated emissions-reduction target relative to 2018 carbon emissions.

Figure 7(b) shows the associated EMTR along the gross income distribution. The optimal carbon policy (red solid line) increases EMTRs relative to the status quo (black dashed line) for incomes above € 13 000. For the highest incomes, this increase is above one percentage point.

Comparison with carbon dividend. We next compare the optimal rebate schedule with a lump-sum carbon dividend policy that achieves the same aggregate emissions reduction. The carbon dividend is illustrated by the blue dotted line in Figure 6. Relative to the optimal policy, the dividend provides larger rebates to tax units in the bottom 45% of the income distribution and smaller rebates to tax units above that threshold. Figure 7(b) shows the corresponding EMTR (blue dotted line). Because the dividend is lump-sum, it does not offset the income gradient in carbon tax liabilities and therefore raises EMTRs relative to the optimal policy. For example, for gross labor income below € 25 000, EMTRs are more than 1 percentage point higher under the carbon dividend. The resulting increase in redistribution is also higher with $\Delta\Pi_{90/10} = 7.40$.[39]

7.2 Along the Abatement Path

We next trace the optimal policy as the emissions cap tightens from the 2018 baseline to zero emissions. For five illustrative levels of emission reduction targets, Table 3 reports the carbon tax rate that implements the corresponding cap.[40] Figure 8(a) shows that the relative incidence of the policy is remarkably stable along this path. Both the regressivity ratio $RR_{90/10}^{\text{rebate}}$ and the progressivity ratio $PR_{90/10}^{\text{carbon tax}}$ vary little with the emission reduction target. The increase in redistribution, however, is non-monotone. As shown in Figure 8(b), the change in the EATR gap $\Delta\Pi_{90/10}$ rises as the cap tightens, reaches its maximum at an emissions reduction of about 65%, and declines thereafter. At the peak, the optimal policy increases the EATR gap by about 4.1 percentage points. A carbon dividend generates a much larger increase of about 15.4 percentage points at the peak.

The same hump shaped pattern is reflected in carbon tax revenue. Figure 8(c) reports revenue in VAT-equivalent terms. The VAT-equivalent increase peaks at nearly 4.5 percentage

[39] See Table 10 in Appendix F for the underlying EATRs.

[40] Figure 18 in Appendix F.1 shows the optimal carbon rebate, EMTR and the net burden of the carbon policy along the gross labor income distribution for an emission reduction target of 10%, 25%, 50%, 75% and 90%.

points when aggregate emissions are reduced by about 60%, before declining as the carbon tax base shrinks.

Finally, Figure 8(d) reports the implied GDP loss along the abatement path, driven by the labor supply reductions due to higher EMTRs. The loss follows a similar hump-shaped profile. Across the abatement path, the GDP loss is substantially larger under a carbon dividend than under the optimal rebate, consistent with the larger increase in EMTRs induced by the dividend policy.

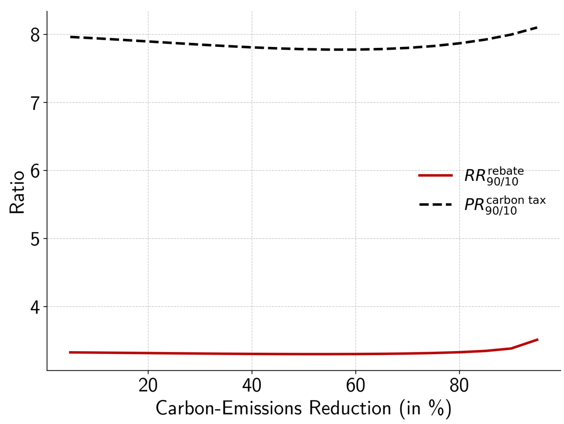

(a) $RR^{\text{rebate}}_{90/10}$ and $PR^{\text{carbon tax}}_{90/10}$

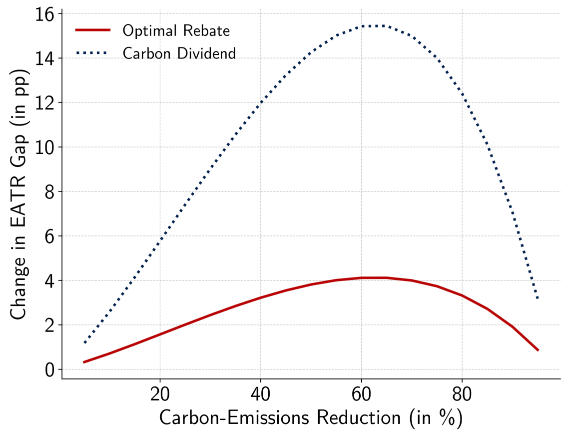

(b) $\Delta\Pi_{90/10}$

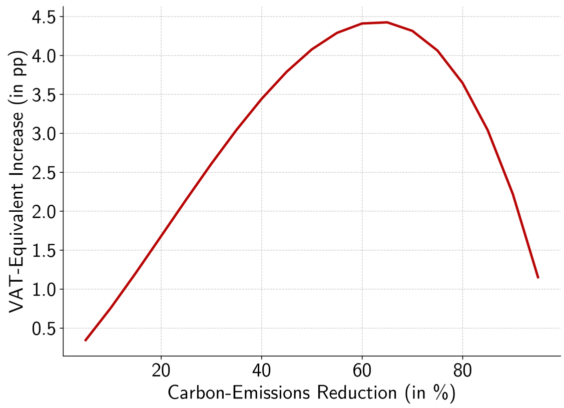

(c) Revenue-equivalent VAT increase

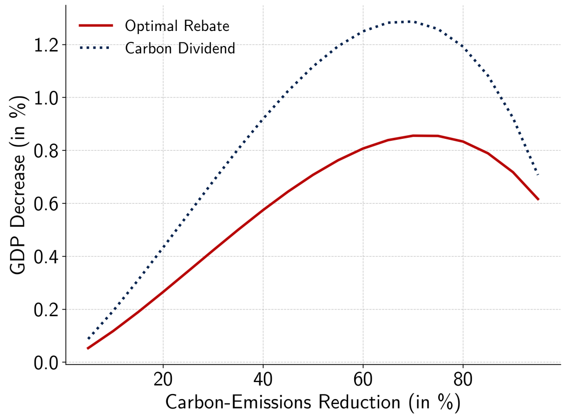

(d) GDP reduction through labor supply responses

Figure 8: Redistributive, Revenue, and GDP Implications of Carbon Policy along the Abatement Path

Notes: The upper left panel shows the regressivity ratio $RR^{\text{rebate}}_{90/10}$ (red solid line) and progressivity ratio $PR^{\text{carbon tax}}_{90/10}$ (black dashed line) for different values of the targeted carbon emission reduction. The upper right panel shows the change in the effective average tax rate (EATR) gap $\Delta\Pi_{90/10}$ for different values of the targeted carbon emission reduction. The red solid line corresponds to the EATR gap under the optimal carbon rebate and the black dashed line to that under a carbon dividend. The lower left panel shows the percentage point revenue-equivalent VAT increase for different values of the targeted carbon emission reduction. The lower right panel shows the percentage reduction in the GDP under the optimal rebate (red solid line) and carbon dividend (black dashed line) for different values of the targeted emission reduction.

Table 4: Robustness of Policy Implications

Case	$RR^{\text{rebate}}_{90/10}$	$PR^{\text{carbon tax}}_{90/10}$	$\Delta\Pi_{90/10}$	τ_{co2}
Baseline	3.31	7.87	2.00	59.91
Lower elasticity ($\varphi = 0.25$)	4.45	7.93	1.22	60.53
Higher elasticity ($\varphi = 0.75$)	2.80	7.84	2.51	59.39
Income-varying elasticity ($\varphi = 0.2 \to \varphi = 0.6$)	4.43	7.92	0.90	59.83
Higher curvature parameter ($\gamma = 1.5$)	4.06	7.92	1.50	60.37
Higher product substitution ($\sigma = 0.5$)	3.49	8.24	1.90	58.02
Lower brown-green substitution ($\sigma_i = 0.5\ \forall i$)	3.33	7.74	2.18	66.41
Lower heating and housing brown-green substitution	3.41	7.81	1.97	62.64
Curvature in welfare function ($\gamma_G = 3$)	4.31	7.92	1.29	60.11
Lower curvature of abatement costs ($\psi = 1$)	3.30	7.77	3.67	109.40
Higher curvature of abatement costs ($\psi = 2$)	3.32	7.90	1.48	44.69

Notes: The table reports distributional patterns of the optimal rebate policy for the baseline specification and ten alternative parameter settings. The case *lower heating and housing brown-green substitution* refers to a scenario in which the elasticity of substitution between the brown and green varieties for housing and heating is set to $\sigma_i = 0.5$, while holding $\sigma_i = 1.5$ for all other products. The case *income-varying elasticity* replaces the constant Frisch elasticity with a smooth logistic profile that equals $\varphi = 0.2$ up to the 10th percentile of the status-quo gross labor income distribution, increases between the 10th and 90th percentiles, and equals $\varphi = 0.6$ above the 90th percentile. The carbon tax τ_{co2} is measured in €/tCO_2e and the change in the EATR gap $\Delta\Pi_{90/10}$ in percentage points (pp). All cases refer to the 25% emission reduction target.

7.3 Robustness

Finally, we assess the robustness of our results to alternative assumptions about labor supply responses, the curvature of consumption utility and the social welfare function, substitution across consumption categories and between brown and green varieties, and the curvature of abatement costs. The results are summarized in Table 4. Appendix F.2 provides a more detailed discussion of all the results.

Overall, our findings are highly robust. Across all specifications, optimal rebates rise with income, but substantially less steeply than carbon tax payments. The progressivity ratio of carbon tax payments is tightly concentrated between 7.75 and 8.25. The rebate ratio varies more, between 2.80 and 4.45, with most of this variation arising from the assumptions about labor-supply elasticities. Thus, the overall carbon policy remains progressive and increases redistribution relative to the status quo in every specification.

Varying the Frisch elasticity mainly affects the rebate schedule. A lower (higher) constant elasticity raises (lowers) the regressivity ratio to 4.45 (2.80) and reduces (increases) the change in the EATR gap to 1.22 (2.51) percentage points. Holding the welfare criterion fixed, a lower elasticity would typically reduce the efficiency cost of redistribution and thereby support more redistribution. In the present robustness exercise, however, the Pareto weights are re-inferred separately for each elasticity specification so that the estimated 2018 German tax–transfer system remains optimal. The lower-elasticity specification is therefore associated with Pareto weights that are less tilted toward lower-income households, counteracting the direct force toward greater redistribution; the reverse holds for the higher elasticity. The income-varying specification resembles the low-elasticity case in terms of the rebate ratio.

Alternative curvatures of abatement costs generate substantial differences in the required carbon price and hence in the scale of the redistributive adjustment, while leaving the relative incidence of rebates and carbon tax payments virtually unchanged. Overall, the alternative specifications affect the magnitude, but not the direction or the central distributional structure, of the optimal policy.

8 Conclusion

In this paper, we show that carbon pricing changes the equity-efficiency trade-off itself. Because lower labor supply also reduces emissions, carbon pricing lowers the welfare cost of redistribution and creates additional scope for redistribution, independently of the social welfare function. This additional scope is captured by the carbon redistribution wedge, the gap between the EMTR and the welfare-relevant labor wedge. A central determinant of this wedge is the marginal propensity to pollute: the increase in emissions generated by an additional unit of expenditure. The higher the marginal propensity to pollute, the more a given reduction in labor supply lowers emissions and, hence, the greater the additional scope for redistribution created by carbon pricing.

Quantifying this mechanism requires measuring both the level and the local slope of carbon Engel curves. We construct consumption-based carbon footprints for German tax units from expenditures on 107 goods and services, combining direct emissions with indirect emissions embodied throughout global production chains. We embed these data in a two-tier non-homothetic CES demand system that captures income-dependent reallocation both across consumption categories and between brown and green varieties within categories, and combine it with endogenous abatement. The resulting framework reproduces the central income gradients in carbon footprints and in the marginal propensity to pollute and permits counterfactual policy analysis under carbon pricing.

In the German application, the optimal rebate at the 90th income percentile is roughly three times that at the 10th percentile, whereas carbon tax payments are almost eight times as large. The rebate therefore rises with income but much less steeply than carbon tax liabilities, so the overall carbon-pricing package is more redistributive than the status quo but less redistributive than a carbon dividend. This central incidence pattern is highly robust across alternative parameter values and emissions-reduction targets.

More broadly, our results show that corrective taxation and redistribution cannot generally be viewed as independent policy problems, even when each objective can be targeted with a separate instrument. In our setting, the carbon tax continues to satisfy the first-best Pigouvian rule, while the income tax remains the instrument for redistribution. Yet carbon pricing changes the welfare cost of labor-supply distortions and therefore the optimal degree of effective redistribution. This mechanism extends beyond carbon emissions. Whenever an externality-generating activity increases with earnings or expenditure and is directly priced, reductions in economic activity induced by redistribution may generate an offsetting external benefit. The EMTR can therefore exceed the welfare-relevant labor wedge.

References

Abrell, J. (2024): "Database for the European Union Transaction Log".

Aghion, P., T. Boppart, M. Peters, M. Schwartzman, and F. Zilibotti (2025): "A theory of endogenous degrowth and environmental sustainability," Discussion paper.

Ahlvik, L., M. Liski, and J. Mäkimattila (2025): "Pigouvian Income Taxation," Discussion paper.

Atkinson, A., and J. Stiglitz (1976): "The design of tax structure: Direct versus indirect taxation," *Journal of Public Economics*, 6(1), 55–75.

Barrage, L. (2020): "Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy," *The Review of Economic Studies*, 87(1), 1–39.

Barrage, L., and W. Nordhaus (2024): "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model," *Proceedings of the National Academy of Sciences*, 121(13), e2312030121.

Berr, N., F. Kindermann, and J. Le Blanc (2026): "From Annual Income to Lifetime Emissions: A Lifecycle Theory of Carbon Inequality," Discussion paper.

Bierbrauer, F. (2025): "Is a Uniform Price on Carbon Desirable?," Discussion paper.

Bigio, S., D. R. Känzig, P. Sánchez, and C. Walsh (2026): "Carbon Pricing and Inequality: A Normative Perspective," *The Economic Journal*.

Blanchard, O. J., C. Gollier, and J. Tirole (2023): "The Portfolio of Economic Policies Needed to Fight Climate Change," *Annual Review of Economics*, 15(1), 693–720.

Blank, P., B. Wickert, A. Obermeier, and R. Friedrich (1999): *Erstellung eines Emissionskatasters für Feuerungsanlagen in Haushalt und Kleinverbrauch: 9 Anlagen*. Umweltbundesamt, https://www.umweltbundesamt.de/publikationen/erstellung-eines-emissionskatasters-fuer-0.

Boskin, M. J., and E. Sheshinski (1978): "Optimal Redistributive Taxation When Individual Welfare Depends Upon Relative Income," *The Quarterly Journal of Economics*, 92(4), 589–601.

Bovenberg, A. L., and R. A. De Mooij (1994): "Environmental levies and distortionary taxation," *The American Economic Review*, 84(4), 1085–1089.

Bovenberg, A. L., and F. van der Ploeg (1994): "Environmental policy, public finance and the labour market in a second-best world," *Journal of Public Economics*, 55(3), 349–390.

Brockjan, K., L. Maier, K. Kott, and N. Sewald (2021): *Umwelt, Energie und Mobilität. Kapitel 13 des Datenreports 2021*. https://www.destatis.de/DE/Service/Statistik-Campus/Datenreport/Downloads/datenreport-2021-kap-13.pdf?__blob=publicationFile.

Bundesagentur für Arbeit (2018): *Ausgaben für aktive und passive Leistungen im SGB II (Jahreszahlen)*.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011): "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *The American Economic Review*, 101(3), 471–475.

Choné, P., and G. Laroque (2010): "Negative Marginal Tax Rates and Heterogeneity," *American Economic Review*, 100(5), 2532–47.

Comin, D., D. Lashkari, and M. Mestieri (2017): "Structural Change With Long-Run Income and Price Effects," Discussion paper.

——— (2021): "Structural Change With Long-Run Income and Price Effects," *Econometrica*, 89(1), 311–374.

Craig, A. C., T. Lloyd, and D. T. Moore (2025): "Do Distributional Concerns Justify Lower Environmental Taxes?," Discussion paper.

Cremer, H., F. Gahvari, and N. Ladoux (1998): "Externalities and Optimal Taxation," *Journal of Public Economics*, 70(3), 343–364.

Dechezleprêtre, A., A. Fabre, T. Kruse, B. Planterose, A. Sanchez Chico, and S. Stantcheva (2025): "Fighting Climate Change: International Attitudes toward Climate Policies," *American Economic Review*, 115(4), 1258–1300.

Douenne, T., S. Dyrda, A. J. Hummel, and M. Pedroni (2025): "Optimal Climate Policy with Incomplete Markets," Discussion paper.

Douenne, T., A. J. Hummel, and M. Pedroni (2026): "Optimal Fiscal Policy in a Climate-Economy Model with Heterogeneous Households," *The Economic Journal*, p. ueag006.

European Environment Agency (2026): *European Union Emissions Trading System (EU ETS) data*. https://www.eea.europa.eu/en/datahub/datahubitem-view/98f04097-26de-4fca-86c4-63834818c0c0.

Familienkasse Direktion SR1 (2018): *Kindergeld / Kinderzuschlag Jahreszahlen 2018.*

Ferey, A., B. B. Lockwood, and D. Taubinsky (2024): "Sufficient statistics for nonlinear tax systems with general across-income heterogeneity," *American Economic Review*, 114(10), 3206–3249.

Ferriere, A., P. Grübener, G. Navarro, and O. Vardishvili (2023): "On the Optimal Design of Transfers and Income Tax Progressivity," *Journal of Political Economy Macroeconomics*, 1(2), 276–333.

Ferriere, A., P. Grübener, and D. Sachs (2025): "Optimal Redistribution: Rising Inequality vs. Rising Living Standards," Discussion paper.

Frank, R. H. (1985): "The Demand for Unobservable and Other Nonpositional Goods," *The American Economic Review*, 75(1), 101–116.

Fricke, L., C. Fuest, and D. Sachs (2025): "Carbon Emissions and Redistribution: The Design of Carbon Tax Rebates," Discussion paper.

German Statistical Office (2022): *Input-Output-Rechnung – Fachserie 18 Reihe 2 (Revision 2019, Stand: August 2022).*

——— (2024a): *VGR des Bundes - Konsumausgaben der privaten Haushalte.*

——— (2024b): *Wohngeld - Ausgaben insgesamt nach Bundesländern im Zeitvergleich am 31.12.*

——— (2025a): *Bruttoausgaben, Einnahmen, Nettoausgaben der Sozialhilfe: Deutschland, Jahre, Sozialhilfearten (22111-0003).*

——— (2025b): *CO2 footprint of private households by areas of demand.* https://www.destatis.de/EN/Themes/Society-Environment/Environment/Environmental-Economic-Accounting/energyflows-emissions/Tables/co2-content-area-of-demand.html.

——— (2025c): *Greenhouse gas emissions (residents concept).* https://www.destatis.de/EN/Themes/Society-Environment/Environment/Environmental-Economic-Accounting/energyflows-emissions/Tables/air-pollutants.html.

Golosov, M., M. Graber, M. Mogstad, and D. Novgorodsky (2023): "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," *The Quarterly Journal of Economics*, 139(2), 1321–1395.

Goulder, L. H. (1995): "Environmental Taxation and the Double Dividend: a Reader's Guide," *International Tax and Public Finance*, 2, 157–183.

Goulder, L. H. (1998): "Environmental Policy Making in a Second-Best Setting," *Journal of Applied Economics*, 1(2), 279–328.

——— (2013): "Climate Change Policy's Interactions with the Tax System," *Energy Economics*, 40, S3–S11, Supplement Issue: Fifth Atlantic Workshop in Energy and Environmental Economics.

Goulder, L. H., M. A. C. Hafstead, G. Kim, and X. Long (2019): "Impacts of a Carbon Tax across U.S. Household Income Groups: What Are the Equity-Efficiency Trade-Offs?," *Journal of Public Economics*, 175, 44–64.

Goulder, L. H., I. W. Parry, R. C. Williams III, and D. Burtraw (1999): "The cost-effectiveness of alternative instruments for environmental protection in a second-best setting," *Journal of Public Economics*, 72(3), 329–360.

Grainger, C. A., and C. D. Kolstad (2010): "Who Pays a Price on Carbon?," *Environmental and Resource Economics*, 46(3), 359–376.

Gruber, J., and E. Saez (2002): "The Elasticity of Taxable Income: Evidence and Implications," *Journal of Public Economics*, 84(1), 1–32.

Hänsel, M. C., M. Franks, M. Kalkuhl, and O. Edenhofer (2022): "Optimal Carbon Taxation and Horizontal Equity: A Welfare-Theoretic Approach with Application to German Household Data," *Journal of Environmental Economics and Management*, 116, 102730.

Hardadi, G., A. Buchholz, and S. Pauliuk (2021): "Implications of the distribution of German household environmental footprints across income groups for integrating environmental and social policy design," *Journal of Industrial Ecology*, 25(1), 95–113.

Heathcote, J., and H. Tsujiyama (2021): "Optimal income taxation: Mirrlees meets Ramsey," *Journal of Political Economy*, 129(11), 3141–3184.

Hendren, N. (2016): "The policy elasticity," *Tax Policy and the Economy*, 30(1), 51–89.

IPCC (2006a): "Chapter 2: Stationary Combustion," in *2006 IPCC Guidelines for National Greenhouse Gas Inventories. Volume 2: Energy*. Institute for Global Environmental Strategies (IGES), for the IPCC.

——— (2006b): "Chapter 3: Mobile Combustion," in *2006 IPCC Guidelines for National Greenhouse Gas Inventories. Volume 2: Energy*. Institute for Global Environmental Strategies (IGES), for the IPCC.

Ireland, N. (2001): "Optimal income tax in the presence of status effects," *Journal of Public Economics*, 81(2), 193–212.

Jacobs, B., and R. A. De Mooij (2015): "Pigou meets Mirrlees: On the irrelevance of tax distortions for the second-best Pigouvian tax," *Journal of Environmental Economics and Management*, 71, 90–108.

Jacobs, B., and F. van der Ploeg (2019): "Redistribution and Pollution Taxes with Non-Linear Engel Curves," *Journal of Environmental Economics and Management*, 95, 198–226.

Jaravel, X., and A. Olivi (2025): "Prices, Non-homotheticities, and Optimal Taxation," Discussion paper.

Kaplow, B. L. (2012): "Optimal control of externalities in the presence of income taxation," *International Economic Review*, 53(2), 487–509.

Kopczuk, W. (2003): "A Note on Optimal Taxation in the Presence of Externalities," *Economics Letters*, 80(1), 81–86.

Levinson, A., and J. O'Brien (2019): "Environmental Engel Curves: Indirect Emissions of Common Air Pollutants," *The Review of Economics and Statistics*, 101(1), 121–133.

Lockwood, B. B. (2020): "Optimal Income Taxation with Present Bias," *American Economic Journal: Economic Policy*, 12(4), 298–327.

Micheletto, L. (2008): "Redistribution and Optimal Mixed Taxation in the Presence of Consumption Externalities," *Journal of Public Economics*, 92(10–11), 2262–2274.

Research Data Centre of the Federal Statistical Office and the statistical offices of the Länder (RDC) (2018a): *Einkommens- und Verbrauchsstichprobe (HB) 2018, On-Site-Zugang*. http://dx.doi.org/10.21242/63221.2018.00.00.1.1.0.

——— (2018b): *Einkommens- und Verbrauchsstichprobe (NGT) 2018, On-Site-Zugang*. http://dx.doi.org/10.21242/63231.2018.00.00.1.1.0.

——— (2018c): *Lohn- und Einkommenssteuerstatistik 2018, Gastwissenschaftsarbeitsplatz*. http://dx.doi.org/10.21242/73111.2018.00.00.2.1.0.

Rothschild, C., and F. Scheuer (2014): "A Theory of Income Taxation under Multidimensional Skill Heterogeneity," Discussion paper.

——— (2016): "Optimal Taxation with Rent-Seeking," *The Review of Economic Studies*, 83(3), 1225–1262.

RWI and forsa (2015): "The German Residential Energy Consumption Survey 2011–2013," Project report, RWI – Leibniz Institute for Economic Research, Essen.

Saez, E. (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *The Review of Economic Studies*, 68(1), 205–229.

——— (2002): "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes," *Journal of Public Economics*, 83(2), 217–230.

Sager, L. (2019): "Income inequality and carbon consumption: Evidence from Environmental Engel curves," *Energy Economics*, 84, 104507.

Sandmo, A. (1975): "Optimal Taxation in the Presence of Externalities," *The Swedish Journal of Economics*, 77(1), 86–98.

Stadler, K., R. Wood, T. Bulavskaya, C.-J. Södersten, M. Simas, S. Schmidt, A. Usubiaga, J. Acosta-Fernández, J. Kuenen, M. Bruckner, S. Giljum, S. Lutter, S. Merciai, J. H. Schmidt, M. C. Theurl, C. Plutzar, T. Kastner, N. Eisenmenger, K.-H. Erb, A. de Koning, and A. Tukker (2018): "EXIOBASE 3: Developing a Time Series of Detailed Environmentally Extended Multi-Regional Input-Output Tables," *Journal of Industrial Ecology*, 22(3), 502–515.

Trading Economics (2026): *EU Carbon Permits: Index Price, Live Quote and Historical Chart*. https://tradingeconomics.com/eecxm:ind.

Umweltbundesamt (2025a): *Carbon Dioxide Emissions for the German Atmospheric Emission Reporting, 1990 - 2024*.

——— (2025b): *Emissionen sinken 2018 um mehr als 31 Prozent gegenüber 1990*. https://www.umweltbundesamt.de/themen/emissionen-sinken-2018-um-mehr-als-31-prozent.

United Nations Climate Change (2026): *Global Warming Potentials (IPCC Fourth Assessment Report)*. https://unfccc.int/process-and-meetings/transparency-and-reporting/greenhouse-gas-data/frequently-asked-questions/global-warming-potentials-ipcc-fourth-assessment-report.

van der Ploeg, F., A. Rezai, and M. Tovar Reaños (2022): "Gathering Support for Green Tax Reform: Evidence from German Household Surveys," *European Economic Review*, 141, 103966.

Weber, C. L., and H. S. Matthews (2008): "Quantifying the global and distributional aspects of American household carbon footprint," *Ecological Economics*, 66(2), 379–391.

Wier, M., M. Lenzen, J. Munksgaard, and S. Smed (2001): "Effects of Household Consumption Patterns on CO2 Requirements," *Economic Systems Research*, 13(3), 259–274.

A Theoretical Appendix

A.1 Household Problem Given Policies and Comparative Statics

This subsection derives the labor supply comparative statics used in Section 3.2 and proves Lemma 1. Throughout this subsection, all objects are evaluated at an arbitrary fixed policy $\mathcal{P}$. To keep the algebra compact, we suppress the argument $\mathcal{P}$ and use the same shorthand for higher derivatives of u as for u_e.

A.1.1 Labor Supply Responses

Implicit differentiation of (7) yields the elasticity of gross income with respect to the retention rate $1-\mathcal{T}'$

$$\varepsilon_{y,1-\mathcal{T}'}(\theta) \equiv \frac{d\log(y(\theta))}{d\log(1-\mathcal{T}')} = \frac{1}{\frac{1}{\varphi}+\Gamma(\theta)\frac{y(\theta)}{e(\theta)}(1-\mathcal{T}'(y(\theta)))+\frac{\mathcal{T}''(y(\theta))y(\theta)}{1-\mathcal{T}'(y(\theta))}},$$

where $\Gamma(\theta) = -\frac{u_{ee}(e(\theta))e(\theta)}{u_e(\theta)}$ is the coefficient of relative risk aversion. Second, the income effect parameter is defined as

$$\eta(\theta) \equiv \frac{\partial y(\theta)}{\partial \mathcal{T}(0)} = \frac{\Gamma(\theta)\frac{y(\theta)}{e(\theta)}}{\frac{1}{\varphi}+\Gamma(\theta)\frac{y(\theta)}{e(\theta)}(1-\mathcal{T}'(y(\theta)))+\frac{\mathcal{T}''(y(\theta))y(\theta)}{1-\mathcal{T}'(y(\theta))}}.$$

For later use, the elasticity of earnings with respect to productivity is

$$\varepsilon_{y,\theta}(\theta) \equiv \frac{d\log y(\theta)}{d\log\theta} = \frac{1+1/\varphi}{\frac{1}{\varphi}+\Gamma(\theta)\frac{y(\theta)}{e(\theta)}(1-\mathcal{T}'(y(\theta)))+\frac{\mathcal{T}''(y(\theta))y(\theta)}{1-\mathcal{T}'(y(\theta))}}.$$

It follows immediately that

$$\frac{\varepsilon_{y,1-\mathcal{T}'}(\theta)}{\varepsilon_{y,\theta}(\theta)} = \frac{\varphi}{1+\varphi}.$$

A.1.2 Proof of Lemma 1

Recall the first-order condition for labor supply (7):

$$u_e(\theta,\mathcal{P})\theta\,(1-\mathcal{T}'(y(\theta))) = \frac{y(\theta)}{\theta}^{\frac{1}{\varphi}}. \tag{36}$$

We prove the result by showing that both tax reforms affect the left-hand side (LHS) in the same way.

Carbon tax reform. We begin from the household's first-stage consumption problem (4). For given expenditure e and carbon tax τ_{co_2}, its Lagrangian is

$$\mathcal{L}^c = U(\mathbf{c}) + \lambda \left[e - \sum_{i=1}^{I} p_i(\tau_{\text{co}_2}) c_i \right].$$

At the optimum, the multiplier on the expenditure constraint equals the marginal utility of expenditure,

$$\lambda(e; \tau_{\text{co}_2}) = u_e(e; \tau_{\text{co}_2}).$$

A marginal increase in the carbon tax affects the first-stage value function through consumer prices. The envelope theorem therefore gives

$$\begin{aligned} u_{\tau_{\text{co}_2}}(e; \tau_{\text{co}_2}) &= -\lambda(e; \tau_{\text{co}_2}) \sum_{i=1}^{I} p_i'(\tau_{\text{co}_2}) c_i(e; \tau_{\text{co}_2}) \\ &= -u_e(e; \tau_{\text{co}_2}) \sum_{i=1}^{I} f_i(\tau_{\text{co}_2}) c_i(e; \tau_{\text{co}_2}) \\ &= -u_e(e; \tau_{\text{co}_2}) f(e; \tau_{\text{co}_2}), \end{aligned} \tag{37}$$

where the second equality uses (3). Thus, at fixed expenditure, a carbon tax increase of $d\tau_{\text{co}_2}$ reduces utility by the marginal utility of expenditure times the household's carbon footprint.

The envelope theorem implies that changes in the household's optimal consumption choices do not enter (37) directly. Similarly, changes in firms' abatement choices are already incorporated in the price response $p_i'(\tau_{\text{co}_2}) = f_i(\tau_{\text{co}_2})$.

To obtain the effect on the marginal utility of expenditure, differentiate (37) with respect to e:

$$\begin{aligned} u_{e\tau_{\text{co}_2}}(e; \tau_{\text{co}_2}) &= \frac{\partial}{\partial e} \left[-u_e(e; \tau_{\text{co}_2}) f(e; \tau_{\text{co}_2}) \right] \\ &= -u_{ee}(e; \tau_{\text{co}_2}) f(e; \tau_{\text{co}_2}) - u_e(e; \tau_{\text{co}_2}) f_e(e; \tau_{\text{co}_2}). \end{aligned} \tag{38}$$

Evaluating this expression at the household's initial expenditure and using $f_e(\theta) = MPP(\theta)$ gives

$$du_e(\theta) = -\left[u_{ee}(\theta) f(\theta) + u_e(\theta) MPP(\theta) \right] d\tau_{\text{co}_2}. \tag{39}$$

The two terms in (39) have distinct interpretations. The term $-u_e(\theta) MPP(\theta) d\tau_{\text{co}_2}$ captures the direct reduction in the utility value of an additional unit of nominal expenditure. The marginal unit of expenditure generates $MPP(\theta)$ units of emissions and therefore becomes more expensive when the carbon tax rises.

The term $-u_{ee}(\theta) f(\theta) d\tau_{\text{co}_2}$ is the real-income component. At fixed nominal expenditure, the carbon tax increase lowers purchasing power by an amount that is, to first order, equivalent to an expenditure loss of $f(\theta)d\tau_{\text{co}_2}$. If $u_{ee} < 0$, this raises the marginal utility of expenditure and partly offsets the first effect.

Hence, the proportional change in the left-hand side of (36) is

$$\frac{du_e(\theta)}{u_e(\theta)} = -\left[MPP(\theta) + \frac{u_{ee}(\theta)}{u_e(\theta)}f(\theta)\right] d\tau_{\text{co}_2}. \tag{40}$$

Income tax reform. Consider the income tax reform

$$d\mathcal{T}\big(y(\theta)\big) = f(\theta)d\tau_{\text{co}_2}.$$

Differentiating this reform along the initial income schedule gives

$$\begin{aligned} d\mathcal{T}'\big(y(\theta)\big) &= \frac{df(\theta)}{dy(\theta)}d\tau_{\text{co}_2} \\ &= MPP(\theta)\left[1 - \mathcal{T}'\big(y(\theta)\big)\right] d\tau_{\text{co}_2}. \end{aligned}$$

At unchanged earnings, the reform reduces expenditure by $f(\theta)d\tau_{\text{co}_2}$ and therefore changes marginal utility by

$$du_e(\theta) = -u_{ee}(\theta)f(\theta)d\tau_{\text{co}_2}.$$

The proportional change in the left-hand side of (36) is consequently

$$\frac{du_e(\theta)}{u_e(\theta)} - \frac{d\mathcal{T}'\big(y(\theta)\big)}{1 - \mathcal{T}'\big(y(\theta)\big)} = -\left[MPP(\theta) + \frac{u_{ee}(\theta)}{u_e(\theta)}f(\theta)\right] d\tau_{\text{co}_2},$$

which coincides with (40). Since the right-hand side of (36) depends on the reform only through the induced change in earnings, the two reforms generate the same first-order labor supply response.

The reforms also have the same first-order utility effect. The carbon tax changes utility by

$$dV(\theta) = u_{\tau_{\text{co}_2}}(\theta)d\tau_{\text{co}_2} = -u_e(\theta)f(\theta)d\tau_{\text{co}_2},$$

while the income tax reform changes utility by

$$dV(\theta) = -u_e(\theta)d\mathcal{T}\big(y(\theta)\big) = -u_e(\theta)f(\theta)d\tau_{\text{co}_2}.$$

The induced earnings adjustments have no first-order utility effect by the envelope theorem.

Finally, the labor supply response to the income tax reform can be decomposed into its substitution and income effects:

$$dy(\theta) = -\frac{\varepsilon_{y,1-\mathcal{T}'}(\theta)y(\theta)}{1 - \mathcal{T}'\big(y(\theta)\big)}\, d\mathcal{T}'\big(y(\theta)\big) + \eta(\theta)d\mathcal{T}\big(y(\theta)\big).$$

Substituting the expressions for $d\mathcal{T}'\big(y(\theta)\big)$ and $d\mathcal{T}\big(y(\theta)\big)$ above yields (9).

A.2 Optimal Income Taxation without a Carbon Cap

This subsection proves Lemma 2 while allowing for general income effects. Throughout the proof, all objects are evaluated at an arbitrary no-cap policy $\mathcal{P} = (\mathcal{T}, 0)$, and we suppress the policy argument to keep the notation compact. At the optimum, $\mathcal{P} = \mathcal{P}_{\text{sq}}$.

The Lagrangian of the government's problem is

$$\mathcal{L} = \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{G}(V(\theta, \mathcal{P}))\, \omega(\theta)\, dH(\theta) + \Lambda \left[\int_{\underline{\theta}}^{\overline{\theta}} \mathcal{T}(y(\theta))\, dH(\theta) - \mathcal{R} \right].$$

Following Saez (2001), fix an interior type θ^* and write $y^* \equiv y(\theta^*)$. Consider a small increase $d\mathcal{T}'$ in the marginal income tax rate over the earnings interval $[y^*, y^* + \Delta y]$. The perturbation is extended continuously so that, up to terms of second order in Δy, every household with earnings above the interval pays an additional $d\mathcal{T}'\Delta y$.

Let $\widetilde{H}$ denote the endogenous earnings distribution. Monotonicity of the earnings schedule implies

$$\widetilde{H}(y(\theta)) = H(\theta), \qquad \widetilde{h}(y(\theta)) = \frac{h(\theta)}{y_\theta(\theta)}.$$

Hence, the mass of households in the affected interval is

$$\widetilde{h}(y^*)\Delta y = \frac{h(\theta^*)\Delta y}{y_\theta(\theta^*)} = \frac{h(\theta^*)\theta^*\Delta y}{\varepsilon_{y,\theta}(\theta^*) y(\theta^*)}.$$

We now collect the substitution, mechanical, and income effects of the perturbation.

Substitution effect. For households in the affected interval, the compensated earnings response is

$$dy^s(\theta^*) = -\varepsilon_{y,1-\mathcal{T}'}(\theta^*) \frac{y(\theta^*)}{1 - \mathcal{T}'(y^*)}\, d\mathcal{T}'.$$

The first-order welfare effect of this response operates through the fiscal externality. Multiplying by the EMTR and by the mass in the interval gives

$$dS(\theta^*) = -\Lambda \frac{\tau_{\text{eff}}(\theta^*)}{1 - \mathcal{T}'(y^*)} \frac{\varepsilon_{y,1-\mathcal{T}'}(\theta^*)}{\varepsilon_{y,\theta}(\theta^*)} \theta^* h(\theta^*)\, d\mathcal{T}'\Delta y. \tag{41}$$

Using the identity derived in Appendix A.1.1,

$$\frac{\varepsilon_{y,1-\mathcal{T}'}(\theta^*)}{\varepsilon_{y,\theta}(\theta^*)} = \frac{\varphi}{1+\varphi},$$

we obtain

$$dS(\theta^*) = -\Lambda \frac{\tau_{\text{eff}}(\theta^*)}{1-\mathcal{T}'(y^*)} \frac{\varphi}{1+\varphi} \theta^* h(\theta^*)\, d\mathcal{T}' \Delta y. \tag{42}$$

Mechanical effect. Households with $\theta > \theta^*$ pay an additional $d\mathcal{T}'\Delta y$. The associated mechanical effect on the Lagrangian is

$$dM(\theta^*) = d\mathcal{T}'\Delta y \int_{\theta^*}^{\overline{\theta}} [\Lambda - g(\theta)]\, dH(\theta). \tag{43}$$

The term Λ is the value of the additional revenue, while $g(\theta)$ is the social value of the household's corresponding loss of expenditure.

Income effect. The increase in tax liability induces an earnings response $\eta(\theta) d\mathcal{T}'\Delta y$ for households above the interval. Its welfare effect is the resulting fiscal externality,

$$dI(\theta^*) = \Lambda d\mathcal{T}'\Delta y \int_{\theta^*}^{\overline{\theta}} \tau_{\text{eff}}(\theta)\eta(\theta) dH(\theta). \tag{44}$$

For later use, a parallel upward shift in the income tax schedule yields the first-order condition

$$0 = -\mathbb{E}[g(\theta)] + \Lambda\left(1 + \mathbb{E}[\tau_{\text{eff}}(\theta)\eta(\theta)]\right).$$

Thus, the shadow value of public funds is

$$\Lambda = \frac{\mathbb{E}[g(\theta)]}{1 + \mathbb{E}[\tau_{\text{eff}}(\theta)\eta(\theta)]}. \tag{45}$$

Optimal tax schedule. Define

$$\begin{aligned} A(\theta^*, \mathcal{P}) &\equiv \frac{\tau_{\text{eff}}(\theta^*, \mathcal{P})}{1 - \mathcal{T}'(y(\theta^*, \mathcal{P})} \frac{\varphi}{1+\varphi} \frac{h(\theta^*)\theta^*}{1 - H(\theta^*)}, \\ B(\mathcal{P}) &\equiv \mathbb{E}[\tau_{\text{eff}}(\theta, \mathcal{P})\eta(\theta, \mathcal{P})], \\ B^+(\theta^*, \mathcal{P}) &\equiv \mathbb{E}[\tau_{\text{eff}}(\theta, \mathcal{P})\eta(\theta, \mathcal{P}) \mid \theta \geq \theta^*]. \end{aligned}$$

At an optimum, $dS(\theta^*) + dM(\theta^*) + dI(\theta^*) = 0$. Dividing this condition by $\Lambda[1 - H(\theta^*)]d\mathcal{T}'\Delta y$ and using (45) gives

$$\begin{aligned} 0 &= -A(\theta^*, \mathcal{P}) + 1 - \frac{\mathbb{E}[g(\theta) \mid \theta \geq \theta^*]}{\Lambda} + B^+(\theta^*, \mathcal{P}) \\ &= -A(\theta^*, \mathcal{P}) + 1 - [1 - D(\theta^*, \mathcal{P})]\,(1 + B(\mathcal{P})) + B^+(\theta^*, \mathcal{P}). \end{aligned}$$

Rearranging and evaluating the condition at $\mathcal{P} = \mathcal{P}_{\text{sq}}$ yields

$$D(\theta^*, \mathcal{P}_{\text{sq}}) = E(\theta^*, \mathcal{P}_{\text{sq}}) \equiv \frac{A(\theta^*, \mathcal{P}_{\text{sq}}) + B(\mathcal{P}_{\text{sq}}) - B^+(\theta^*, \mathcal{P}_{\text{sq}})}{1 + B(\mathcal{P}_{\text{sq}})}. \tag{46}$$

Under Assumption 1, $\eta(\theta, \mathcal{P}_{\text{sq}}) = 0$ for all θ. Moreover, $\tau_{\text{co}_2} = 0$ implies

$$\tau_{\text{eff}}(\theta, \mathcal{P}_{\text{sq}}) = \mathcal{T}'_{\text{sq}}(y(\theta, \mathcal{P}_{\text{sq}}))\,.$$

Equation (46) then reduces to (19), proving Lemma 2.

A.3 Optimal Policies under a Carbon Cap

This subsection proves Lemma 3 and Proposition 2. In contrast to the main text, the proof of Proposition 2 allows for general income effects. The government's Lagrangian is

$$\begin{aligned}
\mathcal{L} = &\int_{\underline{\theta}}^{\overline{\theta}} \mathcal{G}(V(\theta, \mathcal{P}))\,\omega(\theta)\,dH(\theta) \\
&+ \Lambda\left[\int_{\underline{\theta}}^{\overline{\theta}} (\mathcal{T}(y(\theta, \mathcal{P})) + \tau_{\text{co2}} f(\theta, \mathcal{P}))\,dH(\theta) - \mathcal{R}\right] \\
&- \mu\left[\int_{\underline{\theta}}^{\overline{\theta}} f(\theta, \mathcal{P})\,dH(\theta) - \overline{F}\right].
\end{aligned}$$

A.3.1 Proof of Lemma 3

Consider a marginal increase in the carbon tax $d\tau_{\text{co2}}$ accompanied by the opposite of the equivalent income tax reform in (8):

$$d\mathcal{T}(y(\theta, \mathcal{P})) = -f(\theta, \mathcal{P})d\tau_{\text{co2}} \qquad \forall\theta. \tag{47}$$

By Lemma 1, the joint perturbation leaves utility and earnings unchanged to first order. Since $e = y - \mathcal{T}(y)$, the income tax adjustment in (47) raises expenditure at fixed earnings by

$$de(\theta, \mathcal{P}) = f(\theta, \mathcal{P})d\tau_{\text{co2}}.$$

The resulting change in the household's carbon footprint is therefore

$$\begin{aligned}
df(\theta, \mathcal{P}) &= \left[f_{\tau_{\text{co2}}}(\theta, \mathcal{P}) + MPP(\theta, \mathcal{P})f(\theta, \mathcal{P})\right] d\tau_{\text{co2}} \\
&\equiv f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P})d\tau_{\text{co2}}.
\end{aligned} \tag{48}$$

This is the compensated carbon footprint response defined in Section 3.4.

For each household, the direct change in carbon tax payments is

$$d[\tau_{\text{co2}} f(\theta, \mathcal{P})] = f(\theta, \mathcal{P}) d\tau_{\text{co2}} + \tau_{\text{co2}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P}) d\tau_{\text{co2}}.$$

The first term is exactly offset by the income tax reform (47). Hence, the remaining change in government revenue is

$$\tau_{\text{co2}} \int_{\underline{\theta}}^{\overline{\theta}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P})\, dH(\theta) d\tau_{\text{co2}},$$

while aggregate emissions change by

$$\int_{\underline{\theta}}^{\overline{\theta}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P})\, dH(\theta) d\tau_{\text{co2}}.$$

Because household utility is unchanged to first order, the total effect on the Lagrangian is

$$d\mathcal{L} = (\Lambda \tau_{\text{co2}} - \mu) \int_{\underline{\theta}}^{\overline{\theta}} f^c_{\tau_{\text{co2}}}(\theta, \mathcal{P})\, dH(\theta) d\tau_{\text{co2}}. \tag{49}$$

Provided that the aggregate compensated footprint responds locally to the carbon tax, the first-order condition associated with (49) gives

$$\tau_{\text{co2}} = \frac{\mu}{\Lambda},$$

which proves Lemma 3.

A.3.2 Proof of Proposition 2

The proof follows the same local income tax perturbation as Appendix A.2. Throughout this proof, all objects are evaluated at the optimal carbon constrained policy $\mathcal{P}_{\text{co2}}$, and we suppress the policy argument to keep the notation compact. Consider a small increase $d\mathcal{T}'$ in the marginal income tax rate over the earnings interval $[y^*, y^* + \Delta y]$, extended continuously so that households above the interval pay an additional $d\mathcal{T}' \Delta y$.

Relative to the no-cap proof, two changes matter. First, a policy-induced earnings response affects both government revenue and the carbon cap. Its welfare-relevant coefficient is therefore the labor wedge

$$\begin{aligned} \mathrm{T}(\theta) &= \tau_{\text{eff}}(\theta) - \frac{\mu}{\Lambda} \left[1 - \mathcal{T}'(y(\theta))\right] MPP(\theta) \\ &= \tau_{\text{eff}}(\theta) - \Delta_{\text{red}}(\theta) = \mathcal{T}'(y(\theta))\,, \end{aligned} \tag{50}$$

where the last equality follows from Lemma 3. Second, a mechanical change in disposable expenditure changes the carbon footprint even when earnings are fixed. At the optimal carbon tax, however, the associated budget effect $\Lambda\tau_{\text{co2}}df$ and cap effect $-\mu df$ cancel. The mechanical effect of the income tax perturbation is therefore the same as in Appendix A.2.

Substitution effect. The compensated earnings response of households in the affected interval is the same as in Appendix A.2. Multiplying this response by the welfare-relevant labor wedge and by the mass of households in the interval gives

$$\begin{aligned} dS(\theta^*) &= -\Lambda \frac{\mathrm{T}(\theta^*)}{1-\mathcal{T}'(y^*)} \frac{\varepsilon_{y,1-\mathcal{T}'}(\theta^*)}{\varepsilon_{y,\theta}(\theta^*)} \theta^* h(\theta^*)\, d\mathcal{T}'\Delta y \\ &= -\Lambda \frac{\mathrm{T}(\theta^*)}{1-\mathcal{T}'(y^*)} \frac{\varphi}{1+\varphi} \theta^* h(\theta^*)\, d\mathcal{T}'\Delta y, \end{aligned} \tag{51}$$

where the second equality uses the elasticity identity derived in Appendix A.1.1.

Mechanical effect. Households with $\theta > \theta^*$ pay an additional $d\mathcal{T}'\Delta y$. As just explained, the welfare effects of the associated footprint change cancel between the government budget and the carbon cap. Thus,

$$dM(\theta^*) = d\mathcal{T}'\Delta y \int_{\theta^*}^{\overline{\theta}} [\Lambda - g(\theta)]\, dH(\theta). \tag{52}$$

Income effect. The increase in tax liability induces an earnings response $\eta(\theta)d\mathcal{T}'\Delta y$ for households above the interval. Weighting this response by the labor wedge gives

$$dI(\theta^*) = \Lambda d\mathcal{T}'\Delta y \int_{\theta^*}^{\overline{\theta}} \mathrm{T}(\theta)\eta(\theta) dH(\theta). \tag{53}$$

A variation in the lump-sum element of the income tax schedule yields the first-order condition

$$0 = -\mathbb{E}[g(\theta)] + \Lambda\left(1 + \mathbb{E}[\mathrm{T}(\theta)\eta(\theta)]\right).$$

Hence, the shadow value of public funds is

$$\Lambda = \frac{\mathbb{E}[g(\theta)]}{1 + \mathbb{E}[\mathrm{T}(\theta)\eta(\theta)]}. \tag{54}$$

Optimal tax schedule. Define

$$\begin{aligned} A_{\text{co2}}(\theta^*, \mathcal{P}) &\equiv \frac{\tau_{\text{eff}}(\theta^*, \mathcal{P}) - \Delta_{\text{red}}(\theta^*, \mathcal{P})}{1 - \mathcal{T}'_{\text{co2}}(y(\theta^*, \mathcal{P}))} \frac{\varphi}{1+\varphi} \frac{h(\theta^*)\theta^*}{1 - H(\theta^*)}, \\ B_{\text{co2}}(\mathcal{P}) &\equiv \mathbb{E}[\mathrm{T}(\theta, \mathcal{P})\eta(\theta, \mathcal{P})], \\ B^+_{\text{co2}}(\theta^*, \mathcal{P}) &\equiv \mathbb{E}[\mathrm{T}(\theta, \mathcal{P})\eta(\theta, \mathcal{P}) \mid \theta \geq \theta^*]. \end{aligned}$$

At an optimum, $dS(\theta^*)+dM(\theta^*)+dI(\theta^*)=0$. Dividing by $\Lambda[1-H(\theta^*)]d\mathcal{T}'\Delta y$ and using (54) gives

$$\begin{aligned}0 &= -A_{\text{co2}}(\theta^*)+1-\frac{\mathbb{E}[g(\theta)\mid\theta\geq\theta^*]}{\Lambda}+B^+_{\text{co2}}(\theta^*)\\ &= -A_{\text{co2}}(\theta^*)+1-[1-D(\theta^*,\mathcal{P}_{\text{co2}})]\,[1+B_{\text{co2}}(\mathcal{P})]+B^+_{\text{co2}}(\theta^*,\mathcal{P}).\end{aligned}$$

Rearranging yields

$$D(\theta^*,\mathcal{P}_{\text{co2}})=E(\theta^*,\mathcal{P}_{\text{co2}})\equiv\frac{A_{\text{co2}}(\theta^*,\mathcal{P}_{\text{co2}})+B_{\text{co2}}(\mathcal{P}_{\text{co2}})-B^+_{\text{co2}}(\theta^*,\mathcal{P}_{\text{co2}})}{1+B_{\text{co2}}(\mathcal{P}_{\text{co2}})}. \tag{55}$$

Under Assumption 1, $\eta(\theta,\mathcal{P}_{\text{co2}})=0$ for every θ. Equation (55) then reduces to (27), proving Proposition 2.

A.3.3 Proof of Proposition 3

We first examine how a distributionally neutral implementation of the carbon cap affects the distributional gains term. Throughout the proof, all objects without an explicit policy argument are evaluated at $\mathcal{P}_{\text{sq}}$. Under Assumption 2, τ_{co2} is infinitesimal, and we therefore focus throughout on its first-order effects.

Since such a policy leaves utility levels unchanged, only changes in the marginal utility of expenditure can affect the distributional gains term. From the proof of Lemma 1 in Appendix A.1.2, the change in marginal utility induced by a carbon tax increase is given by:[41]

$$\frac{\partial u_e(\theta)}{\partial\tau_{\text{co2}}}d\tau_{\text{co2}}=-\,u_e(\theta)MPP(\theta)\tau_{\text{co2}}-u_{ee}f(\theta)\tau_{\text{co2}}.$$

Next, consider the change in the marginal utility due to the distributionally neutral rebate policy satisfying $R(y(\theta))=\tau_{\text{co2}}f(\theta)$:[42]

$$\frac{\partial u_e(\theta)}{\partial R(y(\theta))}dR(y(\theta))=u_{ee}f(\theta)\tau_{\text{co2}}.$$

Hence, the overall change in marginal utility is given by

$$du_e(\theta)=\frac{\partial u_e(\theta)}{\partial\tau_{\text{co2}}}\tau_{\text{co2}}+\frac{\partial u_e(\theta)}{\partial R(y(\theta))}dR(y(\theta))=-u_e(\theta)MPP(\theta)\tau_{\text{co2}}.$$

[41] Since we consider here an infinitesimal increase of the carbon tax from zero, we have $\tau_{\text{co2}}=d\tau_{\text{co2}}$.
[42] Since we consider here an infinitesimal increase of the carbon rebate from zero, we have $R(y(\theta))=dR(y(\theta))$.

Under Assumption 1, $MPP(\theta) = \alpha\,\forall\theta$, which gives (29).[43] As a consequence, the relative change in marginal utility is the same for all types and therefore we obtain,

$$D(\theta^*, \mathcal{P}_{\text{dn}}) = D(\theta^*, \mathcal{P}_{\text{sq}}). \tag{56}$$

The remainder of the proof, concerning the efficiency costs term, is provided in the main text.

A.3.4 Proof of Proposition 5

We first consider how the distributional gains term is affected by implementing the carbon cap with a carbon dividend. Under Assumption 1, which implies both linear utility in expenditure $u(e)$ and linear carbon Engel curves, the change in marginal utility simplifies to

$$du_e(\theta) = -u_e(\theta)\alpha\tau_{\text{co2}}. \tag{57}$$

This resembles the result from Appendix A.3.3 for the distributionally neutral carbon rebate policy. Yet, under a carbon dividend there is a crucial difference: while the ratio of marginal utilities remains unchanged, the ratio of welfare weights does not, since utility levels shift. Households that receive more in carbon dividends than they pay in carbon taxes experience an increase in utility, whereas households that pay more than they receive experience a decrease in utility. As a result, inequality in utilities falls. To show the implications for the distributional gains term, we substitute the social marginal utility $g(\theta, \mathcal{P})$ into the expression for the distributional gains term:

$$D(\theta^*, \mathcal{P}) = 1 - \frac{E[\mathcal{G}'\left(V(\theta, \mathcal{P})\right)\omega(\theta)u_e(\theta, \mathcal{P})|\theta \geq \theta^*]}{E[\mathcal{G}'\left(V(\theta, \mathcal{P})\right)\omega(\theta)u_e(\theta, \mathcal{P})]}. \tag{58}$$

Formally, the carbon dividend affects the utility level of type θ by

$$dV(\theta) = u_e(\theta)\left(\bar{R}_{div} - \tau_{\text{co2}}\alpha e(\theta)\right),$$

implying that $\mathcal{G}'(V(\theta, \mathcal{P}))$ changes by

$$d\mathcal{G}'\left(V(\theta, \mathcal{P})\right) = \mathcal{G}''\left(V(\theta, \mathcal{P})\right)u_e(\theta)\left(\bar{R}_{div} - \tau_{\text{co2}}\alpha e(\theta)\right).$$

Hence, the relative change of the numerator of (58) (also accounting for (57)), is given by:

$$\frac{E[\mathcal{G}''\left(V(\theta, \mathcal{P})\right)\omega(\theta)u_e(\theta)^2\left(\bar{R}_{div} - \tau_{\text{co2}}\alpha e(\theta)\right) - \mathcal{G}'\left(V(\theta, \mathcal{P})\right)\omega(\theta)u_e(\theta)\alpha\tau_{\text{co2}}|\theta \geq \theta^*]}{E[\mathcal{G}'\left(V(\theta, \mathcal{P})\right)\omega(\theta)u_e(\theta, \mathcal{P})|\theta \geq \theta^*]}. \tag{59}$$

[43]Note that due to the linearity of $u(e; \mathcal{P})$ in e, we could have obtained this results with fewer steps since $u_{ee} = 0$. Yet, we kept it more general because this result about the distributional gains term holds also for utility functions with curvature in e.

Analogously, for the denominator, it is

$$\frac{E[\mathcal{G}''\left(V(\theta,\mathcal{P})\right)\omega(\theta)u_e(\theta)^2\left(\bar{R}_{div}-\tau_{\text{co2}}\alpha e(\theta)\right)-\mathcal{G}'\left(V(\theta,\mathcal{P})\right)\omega(\theta)u_e(\theta)\alpha\tau_{\text{co2}}]}{E[\mathcal{G}'\left(V(\theta,\mathcal{P})\right)\omega(\theta)u_e(\theta,\mathcal{P})]}.$$

We now define the weighted expectation operator

$$E_g[X(\theta)\mid\theta\geq x]\equiv\frac{E[g(\theta)X(\theta)\mid\theta\geq x]}{E[g(\theta)\mid\theta\geq x]}.$$

Hence, we can write (59) as

$$E_g\Big[\frac{\mathcal{G}''}{\mathcal{G}'}V\frac{u_e}{V}\left(\bar{R}_{div}-\tau_{\text{co2}}\alpha e(\theta)\right)-\alpha\tau_{\text{co2}}|\theta\geq x\Big],$$

The second part in this expectation operator is independent of θ^*. We therefore define

$$\xi(x)=E_g\Big[\Psi(\theta)|\theta\geq x\Big],$$

where

$$\Psi(\theta)=\underbrace{-\frac{\mathcal{G}''}{\mathcal{G}'}V}_{CRUIA(V(\theta))}\frac{u_e}{V}\left(\tau_{\text{co2}}\alpha e(\theta)-\bar{R}_{div}\right).$$

The proportional changes in the numerator and denominator of (58) are therefore $\xi(\theta^*)-\alpha\tau_{\text{co2}}$ and $\xi(\underline{\theta})-\alpha\tau_{\text{co2}}$, respectively. Differentiating the ratio in (58) gives

$$dD(\theta^*)=-\big[1-D(\theta^*,\mathcal{P}_{\text{sq}})\big]\left[\xi(\theta^*)-\xi(\underline{\theta})\right].$$

Since $1-D(\theta^*,\mathcal{P}_{\text{sq}})>0$, a necessary and sufficient condition for the distributional gains term to decrease due to the carbon dividend is

$$\xi(\theta^*)>\xi(\underline{\theta}). \tag{60}$$

There exists a threshold $\tilde{\theta}$ such that $\Psi(\theta)>(<)0$ if $\theta>(<)\tilde{\theta}$. Intuitively, $\tilde{\theta}$ is the productivity type for whom the carbon dividend equals exactly the carbon tax burden. We now show that $\Psi'(\theta)>0$ for each θ which then implies (60).

To establish this result, write the affine status-quo income tax schedule as $\mathcal{T}_{\text{sq}}(y)=ty-b$, where $t=\mathcal{T}'_{\text{sq}}$ and $b=-\mathcal{T}_{\text{sq}}(0)$. Hence, $e(\theta)=(1-t)y(\theta)+b$. Under Assumption 1, indirect utility from consumption is linear in expenditure, so total indirect utility under the status quo is

$$V(\theta)=u_ee(\theta)-\frac{(y(\theta)/\theta)^{1+1/\varphi}}{1+1/\varphi}.$$

The labor supply first-order condition is

$$u_e\theta(1-t)=\left(\frac{y(\theta)}{\theta}\right)^{1/\varphi},$$

which can be shown to imply

$$V(\theta)=u_e\big[(1-t)y(\theta)+b\big]-\frac{\varphi}{1+\varphi}u_e(1-t)y(\theta)=\frac{u_e}{1+\varphi}\big[e(\theta)+\varphi b\big].$$

The proportional utility change induced by the carbon dividend is therefore

$$\frac{dV(\theta)}{V(\theta)}=(1+\varphi)\frac{\bar{R}_{div}-\tau_{\text{co2}}\alpha e(\theta)}{e(\theta)+\varphi b}. \tag{61}$$

Differentiating with respect to θ gives

$$\frac{d}{d\theta}\left(\frac{dV(\theta)}{V(\theta)}\right)=-(1+\varphi)\frac{\bar{R}_{div}+\varphi b\tau_{\text{co2}}\alpha}{\left[e(\theta)+\varphi b\right]^2}e'(\theta)<0, \tag{62}$$

since $e'(\theta)>0$, $\bar{R}_{div}>0$ and $\varphi b\tau_{\text{co2}}\alpha>0$.

Under the isoelastic welfare function assumed in Proposition 5,

$$-\frac{\mathcal{G}''(V)V}{\mathcal{G}'(V)}=\gamma_G.$$

The type-specific component of the proportional change in social marginal utility can therefore be written as

$$\Psi(\theta)=-\gamma_G\frac{dV(\theta)}{V(\theta)}.$$

Equation (62) then implies

$$\Psi'(\theta)>0 \qquad \forall\theta.$$

Because $g(\theta)$ defines positive weights, the monotonicity of $\Psi(\theta)$ implies

$$E_g\left[\Psi(\theta)\mid\theta\geq\theta^*\right]>E_g\left[\Psi(\theta)\mid\theta<\theta^*\right] \quad \forall\theta^*\in]\underline{\theta},\overline{\theta}[.$$

The unconditional weighted mean $\xi(\underline{\theta})$ is a weighted average of the means below and above θ^*. Consequently,

$$\xi(\theta^*)>\xi(\underline{\theta}) \qquad \forall\theta^*\in]\underline{\theta},\overline{\theta}[.$$

By condition (60), this establishes

$$D(\theta^*,\mathcal{P}_{\text{div}})<D(\theta^*,\mathcal{P}_{\text{sq}}) \qquad \forall\theta^*\in]\underline{\theta},\overline{\theta}[.$$

Finally, a carbon dividend changes the income tax schedule only through its lump-sum component. Together with Assumption 3, this implies

$$E(\theta^*, \mathcal{P}_{\text{div}}) = E(\theta^*, \mathcal{P}_{\text{sq}}) \qquad \forall \theta^*.$$

Using the status-quo optimality condition, we therefore obtain

$$E(\theta^*, \mathcal{P}_{\text{div}}) = E(\theta^*, \mathcal{P}_{\text{sq}}) = D(\theta^*, \mathcal{P}_{\text{sq}}) > D(\theta^*, \mathcal{P}_{\text{div}})$$

for every interior θ^*, which proves Proposition 5.

Implications of decreasing CRUIA. We finally discuss the role of the assumption that relative utility-inequality aversion is constant. Let

$$C(V) \equiv -\frac{\mathcal{G}''(V)V}{\mathcal{G}'(V)}$$

and denote the proportional utility change by

$$r(\theta) \equiv \frac{dV(\theta)}{V(\theta)}.$$

Then

$$\Psi(\theta) = -C(V(\theta))r(\theta)$$

and

$$\Psi'(\theta) = -C'(V(\theta))V'(\theta)r(\theta) - C(V(\theta))r'(\theta). \tag{63}$$

We have shown above that $r'(\theta) < 0$. Moreover, $r(\theta) > 0$ for $\theta < \tilde{\theta}$ and $r(\theta) < 0$ for $\theta > \tilde{\theta}$.

Suppose now that relative utility-inequality aversion is decreasing, $C'(V) \leq 0$. For $\theta < \tilde{\theta}$, both terms on the right-hand side of (63) are non-negative, and the second is strictly positive. Hence,

$$\Psi'(\theta) > 0 \qquad \forall \theta < \tilde{\theta}.$$

Therefore, for any cutoff $\theta^* < \tilde{\theta}$, every value of $\Psi(\theta)$ above the cutoff exceeds every value below it. Between θ^* and $\tilde{\theta}$, this follows from the monotonicity of Ψ; above $\tilde{\theta}$, it follows from $\Psi(\theta) > 0$, whereas $\Psi(\theta) < 0$ below the cutoff. Thus,

$$\xi(\theta^*) > \xi(\underline{\theta}) \qquad \forall \theta^* < \tilde{\theta}.$$

Decreasing relative utility-inequality aversion therefore does not change the direction of the local reform for cutoffs below the type that is neither a net beneficiary nor a net contributor.

For $\theta > \tilde{\theta}$, by contrast, the first term in (63) is non-positive, while the second term remains positive. If relative utility-inequality aversion decreases sufficiently strongly, the first term can dominate and $\Psi'(\theta)$ can become negative. For cutoffs $\theta^* > \tilde{\theta}$, condition (60) may then fail. At sufficiently high incomes, the welfare improving local reform may therefore point toward greater redistribution than under the carbon dividend, corresponding to a locally decreasing rebate. Economically, this corresponds to additional redistribution within the upper tail of the income distribution: although households at the very top experience the larger proportional utility loss, sufficiently decreasing relative utility-inequality aversion can make the social marginal utility of households somewhat lower in the upper tail increase proportionally more than that of households at the very top. Compensating the latter for the additional burden from the carbon dividend then requires imposing an even larger burden on households at the very top.

B Computation of Emission Intensities

We compute emission intensities by combining two data sources: (i) environmentally extended multi-regional input-output (EE-MRIO) data from EXIOBASE v3.9.6 (Stadler et al. 2018) and (ii) German supply-use and input-output tables (SUT/IOT) from the national accounts provided by the German Federal Statistical Office (German Statistical Office 2022). EXIOBASE provides the global input-output framework needed to compute emission intensities at basic prices, but it does not provide supply-use tables required to convert them to purchaser prices. We therefore use the German SUT/IOT to recover product-level trade margins and net product taxes, which we then map into the EXIOBASE framework.

Section B.1 describes the German SUT/IOT data and Section B.2 describes the EXIOBASE data and explains how we map the German trade margins and net product taxes into the EXIOBASE framework. Section B.3 then describes the computation of emission intensities and Section B.4 validates these intensities by comparing the aggregate emissions implied by our approach with official estimates reported by German Statistical Office (2025b).

B.1 German Supply-Use and Input-Output Tables

B.1.1 Data

The *supply table* of the German SUT reports total supply at both basic and purchaser prices, together with trade margins and taxes on products less subsidies.[44] For household final consumption, the *use table* is reported only at purchaser prices and therefore does not separately identify basic prices, trade margins, and net product taxes. We recover this decomposition by combining the SUT with the IOT, which report household final consumption expenditure at

[44] By construction, supply at purchaser prices equals supply at basic prices plus trade margins plus net taxes on products.

basic prices. The SUT and IOT are published in the CPA classification, but at different levels of disaggregation: the SUT distinguish 85 CPA products, whereas the IOT distinguish 72 CPA products. We work at the 85-product SUT level.[45]

In the following, we recover household final consumption at basic prices, trade margins, and net product taxes in the 85-product CPA classification. In Section B.2, we then map these components to the ISIC classification used in EXIOBASE, thereby obtaining an EXIOBASE-consistent use table.

B.1.2 Construction of German Use Table

We recover household final consumption at basic prices, trade margins, and net product taxes in four steps. We first harmonize the 72-product IOT and 85-product SUT classifications. We then recover the total wedge between purchaser and basic prices for household consumption, decompose this wedge into trade margins and net product taxes, and finally balance trade margins so that they satisfy the accounting identity that they sum to zero.

Harmonizing IOT and SUT classifications. Household final consumption expenditure at basic prices is available only in the IOT, which distinguish 72 products, whereas the SUT distinguish 85 products. To reconcile the two classifications, we define a binary aggregation matrix $\mathbf{S} \in \{0,1\}^{72\times 85}$, where $S_{ik} = 1$ if product k in the 85 product classification belongs to product group i in the 72-product classification, and $S_{ik} = 0$ otherwise. Each 85-product category maps to exactly one 72-product category. For each 72-product category c, let $\mathcal{I}_c := \{k \in \{1,\dots,85\} : S_{ck} = 1\}$ denote the set of corresponding 85-product categories. The aggregation matrix $\mathbf{S}$ aggregates any 85-product vector to the 72-product classification by summing within categories.[46]

Recovering the wedge between purchaser and basic prices. Let $\mathbf{y}_{pp}^{SUT,85} \in \mathbb{R}^{85\times 1}$ denote household final consumption expenditure at purchaser prices from the SUT *use table*, and let $\mathbf{y}_{bp}^{IOT,72} \in \mathbb{R}^{72\times 1}$ denote household final consumption expenditure at basic prices from the IOT. Aggregating the SUT use table to the 72-product classification gives

$$\mathbf{y}_{pp}^{SUT,72} = \mathbf{S}\mathbf{y}_{pp}^{SUT,85}.$$

[45] An alternative would be to work at the 72-product IOT classification and aggregate the SUT from 85 to 72 CPA products. We do not perform the conversion from purchaser to basic prices at that level because aggregation to 72 products removes within-group heterogeneity in the wedge between purchaser and basic prices. In particular, products that belong to the same 72-product category would inherit the same average wedge, even when trade margins and net product taxes differ substantially across the underlying 85-product categories. Working at the 85-product level preserves this variation. Our results are robust using this alternative approach.

[46] Formally, aggregation is given by $\mathbf{x}^{72} = \mathbf{S}\mathbf{x}^{85}$. Conversely, a 72-product vector can be disaggregated to 85 products using weights $\mathbf{w}^{85} \in \mathbb{R}_{+}^{85\times 1}$ satisfying $\mathbf{S}\mathbf{w}^{85} = \mathbf{1}_{72}$, so that $\mathbf{x}^{85} = (\mathbf{S}^{\intercal}\mathbf{x}^{72}) \odot \mathbf{w}^{85}$, where $\odot$ denotes the Hadamard product.

The implied wedge between purchaser and basic prices in household consumption at the 72-product level is then

$$\boldsymbol{\delta}^{SUT,72} = \mathbf{y}_{pp}^{SUT,72} - \mathbf{y}_{bp}^{IOT,72}.$$

This wedge reflects both trade margins and net product taxes.

We disaggregate $\boldsymbol{\delta}^{SUT,72}$ to the 85-product classification using information from the SUT *supply table*. Let $\mathbf{s}_{pp}^{SUT,85}$ and $\mathbf{s}_{bp}^{SUT,85}$ denote total supply at purchaser and basic prices, and define the corresponding supply-side wedge as $\boldsymbol{\delta}_{supply}^{SUT,85} = \mathbf{s}_{pp}^{SUT,85} - \mathbf{s}_{bp}^{SUT,85}$. We assume that, within each 72-product category, the household wedge is distributed across 85-product categories in proportion to the supply-side wedge. Accordingly, for each $k \in \mathcal{I}_c$

$$w_{\delta,k}^{SUT,85} = \frac{\delta_{supply,k}^{SUT,85}}{\sum_{j\in\mathcal{I}_c} \delta_{supply,j}^{SUT,85}}.$$

By construction, these weights sum to one within each 72-product category.[47] The disaggregated household wedge is given by

$$\boldsymbol{\delta}^{SUT,85} = (\mathbf{S}^{\intercal}\boldsymbol{\delta}^{SUT,72}) \odot \mathbf{w}_{\boldsymbol{\delta}}^{SUT,85},$$

where $\odot$ denotes the Hadamard product.

To ensure nonnegative household consumption at basic prices, we impose the admissibility constraint $\delta_k^{SUT,85} \leq y_{pp,k}^{SUT,85}$. Whenever the initial allocation violates this bound for some product k, we set $\delta_k^{SUT,85} = y_{pp,k}^{SUT,85}$, and reallocate the residual wedge within the same 72-product group across the remaining products in proportion to their renormalized weights.[48]

We then recover household final consumption expenditure at basic prices in the 85-product SUT classification as

$$\mathbf{y}_{bp}^{SUT,85} = \mathbf{y}_{pp}^{SUT,85} - \boldsymbol{\delta}^{SUT,85}.$$

This procedure recovers household consumption at basic prices in the 85-product classification while remaining consistent with the published SUT and IOT totals. Finally, we define the corresponding basic-to-purchaser price conversion factor as

$$\tilde{\boldsymbol{\delta}}^{SUT,85} = \mathbf{y}_{pp}^{SUT,85} \oslash \mathbf{y}_{bp}^{SUT,85},$$

where $\oslash$ denotes element-wise division. Each element $\tilde{\delta}_k^{SUT,85}$ converts household consumption expenditure from basic to purchaser prices for product k.

[47] Whenever $\sum_{j\in\mathcal{I}_c} \delta_{supply,j}^{SUT,85} = 0$, we set $w_{\delta,k}^{SUT,85} = 0$ for all $k \in \mathcal{I}_c$. In this case the weights sum to zero within each 72-product category. In our data, this occurs only when the corresponding household wedge $\delta_c^{SUT,72}$ is also zero, meaning the group-level accounting identity is preserved.

[48] In the data, this adjustment affects only two products, and the reallocated amount is quantitatively small.

Trade margins and net product taxes. We next decompose the household wedge between purchaser and basic prices $\boldsymbol{\delta}^{SUT,85}$ into trade margins and net product taxes using the decomposition of the corresponding wedge from the SUT *supply table*. Let $\mathbf{m}^{SUT,85}_{supply}$ denote trade margins and $\mathbf{t}^{SUT,85}_{supply}$ denote net product taxes for total supply at the 85-product level, so that $\boldsymbol{\delta}^{SUT,85}_{supply} = \mathbf{m}^{SUT,85}_{supply} + \mathbf{t}^{SUT,85}_{supply}$. For each product $k = 1, \ldots, 85$ with $\delta^{SUT,85}_{supply,k} \neq 0$, define the supply-side decomposition coefficients

$$\tilde{m}^{SUT,85}_{supply,k} = \frac{m^{SUT,85}_{supply,k}}{\delta^{SUT,85}_{supply,k}}, \qquad \tilde{t}^{SUT,85}_{supply,k} = \frac{t^{SUT,85}_{supply,k}}{\delta^{SUT,85}_{supply,k}}.$$

By construction, $\tilde{m}^{SUT,85}_{supply,k} + \tilde{t}^{SUT,85}_{supply,k} = 1$. If $\delta^{SUT,85}_{supply,k} = 0$, we set $\tilde{m}^{SUT,85}_{supply,k} = \tilde{t}^{SUT,85}_{supply,k} = 0$.[49]

We then assume that the wedge from the *use table* is split between trade margins and net product taxes in the same way as the supply-side wedge. Household trade margins and household net product taxes are therefore given by

$$\mathbf{m}^{SUT,85} = \tilde{\boldsymbol{m}}^{SUT,85}_{supply} \odot \boldsymbol{\delta}^{SUT,85}, \qquad \mathbf{t}^{SUT,85} = \tilde{\boldsymbol{t}}^{SUT,85}_{supply} \odot \boldsymbol{\delta}^{SUT,85}.$$

By construction, $\mathbf{m}^{SUT,85} + \mathbf{t}^{SUT,85} = \boldsymbol{\delta}^{SUT,85}$, so that $\mathbf{y}^{SUT,85}_{bp} = \mathbf{y}^{SUT,85}_{pp} - \mathbf{m}^{SUT,85} - \mathbf{t}^{SUT,85}$.

Balancing trade margins. The decomposition of the wedge between purchaser and basic prices need not satisfy the accounting identity that trade margins sum to zero in aggregate. We therefore rebalance trade margins and net product taxes to impose that restriction while preserving the wedge between purchaser and basic prices $\boldsymbol{\delta}^{SUT,85}$.

Let $P := \{k : m^{SUT,85}_k > 0\}$ and $N := \{k : m^{SUT,85}_k < 0\}$ denote the sets of products with positive and negative preliminary trade margins. Define the scaling factor

$$\lambda^{SUT,85} \equiv \frac{\sum_{k \in N} |m^{SUT,85}_k|}{\sum_{k \in P} m^{SUT,85}_k}.$$

We then rescale positive trade margins proportionally and leave negative trade margins unchanged:

$$\hat{m}^{SUT,85}_k = \begin{cases} \lambda^{SUT,85} m^{SUT,85}_k, & k \in P, \\ m^{SUT,85}_k, & \text{otherwise.} \end{cases}$$

The amount removed from positive trade margins is added to net product taxes:

$$\hat{t}^{SUT,85}_k = t^{SUT,85}_k + \left(m^{SUT,85}_k - \hat{m}^{SUT,85}_k\right).$$

[49] Whenever $y^{SUT,85}_{pp,k} = 0$ *and* $\tilde{t}^{SUT,85}_{supply,k} > 0$, we set the tax component to zero and assign the entire wedge between basic and purchaser prices to trade margins, i.e. $\tilde{m}^{SUT,85}_{supply,k} = 1$ and $\tilde{t}^{SUT,85}_{supply,k} = 0$. This prevents the allocation of net product taxes to products with zero household final consumption expenditure at purchaser prices while preserving the accounting identity $\tilde{m}^{SUT,85}_{supply,k} + \tilde{t}^{SUT,85}_{supply,k} = 1$. In the data, this adjustment affects only two products and is quantitatively negligible.

By construction, $\hat{m}_k^{SUT,85} + \hat{t}_k^{SUT,85} = \delta_k^{SUT,85}$ for every product k and $\sum_{k=1}^{85} \hat{m}_k^{SUT,85} = 0$.

B.2 EXIOBASE

B.2.1 Data

EXIOBASE provides monetary input-output tables together with physical extensions, including detailed emission accounts. The air emissions extension covers 32 distinct pollutants, disaggregated by products and final demand. EXIOBASE covers 44 individual countries plus five 'rest of the world' regions giving a total of $K = 49$ regions. We use the product-by-product dimension of EXIOBASE, which distinguishes $P = 200$ products, giving $K \times P = 49 \times 200 = 9,800$ region-product combinations. The tables are provided in basic prices, with products classified according to the EXIOBASE product classification (Stadler et al. 2018).

B.2.2 Construction of EXIOBASE Use Table

We use the German household final-consumption decomposition constructed in Section B.1.2 to convert household final consumption expenditure in EXIOBASE from basic to purchaser prices. Specifically, we map the basic-to-purchaser price conversion factors, as well as the decomposition of the purchaser-basic price wedge into trade margins and net product taxes, from the CPA classification to the EXIOBASE product classification. Applying these mapped objects to EXIOBASE household final consumption at basic prices yields a purchaser-price household final-consumption vector together with its decomposition into basic prices, trade margins, and net product taxes.

Correspondence from CPA to EXIOBASE. To map the basic-to-purchaser price conversion factor and the decomposition coefficients from the CPA classification to the EXIOBASE product classification, we use a correspondence matrix provided by Stadler et al. (2018) .[50] Let $\mathbf{C} \in \{0,1\}^{200\times 85}$ denote the binary correspondence matrix, where $C_{pk} = 1$ if CPA product k is mapped to EXIOBASE product p, and $C_{pk} = 0$ otherwise. Because several CPA products may map to the same EXIOBASE product, we aggregate the CPA-level basic-to-purchaser price conversion factors and decomposition coefficients using household final consumption expenditure at basic prices, $\mathbf{y}_{bp}^{SUT,85}$, as weights. We therefore compute each EXIOBASE-level value as a weighted average over the CPA products associated with that EXIOBASE product. For any

[50] Strictly speaking, Stadler et al. (2018) provides a NACE-to-EXIOBASE correspondence matrix. Since each CPA product can be assigned to a single NACE category, the same concordance can be applied to CPA products.

CPA-level vector $\mathbf{x}^{SUT,85} \in \mathbb{R}^{85\times 1}$, the corresponding vector in the EXIOBASE classification $\mathbf{x}^{EXIO,200} \in \mathbb{R}^{200\times 1}$ is given by:

$$x_p^{EXIO,200} = \frac{\sum_{k=1}^{85} C_{pk}\, y_{bp,k}^{SUT,85}\, x_k^{85}}{\sum_{k=1}^{85} C_{pk}\, y_{bp,k}^{SUT,85}} \qquad \forall\, p = 1, \ldots, 200.$$

Household final consumption expenditure at purchaser prices. We first map the basic-to-purchaser price conversion factor $\tilde{\boldsymbol{\delta}}^{SUT,85}$ from the CPA classification to the EXIOBASE product classification, yielding $\tilde{\boldsymbol{\delta}}^{EXIO,200} \in \mathbb{R}^{200\times 1}$. Applying this factor to household final consumption expenditure at basic prices in EXIOBASE $\mathbf{y}_{bp}^{EXIO,200} \in \mathbb{R}^{200\times 1}$, gives household final consumption expenditures at purchaser prices:

$$\mathbf{y}_{pp}^{EXIO,200} = \mathbf{y}_{bp}^{EXIO,200} \odot \tilde{\boldsymbol{\delta}}^{EXIO,200}.$$

The corresponding wedge between basic and purchaser prices is then given by

$$\boldsymbol{\delta}^{EXIO,200} = \mathbf{y}_{pp}^{EXIO,200} - \mathbf{y}_{bp}^{EXIO,200}.$$

Trade margins and net product taxes. Finally, we decompose the wedge between purchaser and basic prices $\boldsymbol{\delta}^{EXIO,200}$ into trade margins and net product taxes. Starting from the CPA-level decomposition $\hat{\mathbf{m}}^{SUT,85}$ and $\hat{\mathbf{t}}^{SUT,85}$, we define

$$r_{m,k}^{SUT,85} = \frac{\hat{m}_k^{SUT,85}}{|\delta_k^{SUT,85}|}, \qquad r_{t,k}^{SUT,85} = \frac{\hat{t}_k^{SUT,85}}{|\delta_k^{SUT,85}|},$$

for $|\delta_k^{SUT,85}| > 0$, and set them to zero otherwise. We normalize by the absolute wedge because some products carry negative trade margins, reflecting their role as suppliers of trade services. These normalized components therefore preserve the sign of each margin while scaling by the magnitude of the purchaser-basic price wedge. By construction,

$$r_{m,k}^{SUT,85} + r_{t,k}^{SUT,85} = \operatorname{sgn}(\delta_k^{SUT,85}).$$

We then map $\mathbf{r}_m^{SUT,85}$ and $\mathbf{r}_t^{SUT,85}$ to the EXIOBASE product classification using the correspondence described above, yielding $\mathbf{r}_m^{EXIO,200} \in \mathbb{R}^{200\times 1}$ and $\mathbf{r}_t^{EXIO,200} \in \mathbb{R}^{200\times 1}$. Let $\mathbf{s}^{EXIO,200} \equiv \mathbf{r}_m^{EXIO,200} + \mathbf{r}_t^{EXIO,200}$. We recover EXIOBASE trade margins and net product taxes as

$$m_p^{EXIO,200} = \frac{r_{m,p}^{EXIO,200}}{s_p^{EXIO,200}}\, \delta_p^{EXIO,200}, \qquad t_p^{EXIO,200} = \frac{r_{t,p}^{EXIO,200}}{s_p^{EXIO,200}}\, \delta_p^{EXIO,200}.$$

Whenever $s_p^{EXIO,200} = 0$, we set $m_p^{EXIO,200} = t_p^{EXIO,200} = 0$. This yields a decomposition of the purchaser-basic price wedge into trade margins and net product taxes.

Balancing trade margins. As in the German *use table*, the decomposition of the purchaser-basic price wedge in the EXIOBASE product classification need not satisfy the accounting identity that trade margins sum to zero in the aggregate. We therefore apply the balancing procedure described in Section B.1.2 to the EXIOBASE vectors $\mathbf{m}^{EXIO,200}$ and $\mathbf{t}^{EXIO,200}$. Positive trade margins are scaled down proportionally so that they exactly offset negative trade margins, and the removed amount is added product by product to net product taxes. We denote the balanced tax vector and trade margin by $\hat{\mathbf{t}}^{EXIO,200}$ and $\hat{\mathbf{m}}^{EXIO,200}$, respectively. By construction, $\hat{\mathbf{m}}^{EXIO,200} + \hat{\mathbf{t}}^{EXIO,200} = \boldsymbol{\delta}^{EXIO,200}$ and $\mathbf{1}_{200}^{\top}\hat{\mathbf{m}}^{EXIO,200} = 0$.

B.3 Computation of Emission Intensities

In this section we describe the details of the computation of indirect and direct carbon emission intensities.

Carbon dioxide equivalents. In line with the Kyoto Protocol, we measure greenhouse gas emissions in carbon dioxide equivalents (CO_2e).[51] Emissions of carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), sulfur hexafluoride (SF_6), Hydrofluorocarbons (HFC), and Perfluorocarbons (PFC) are converted into CO_2e using 100-year global warming potentials (GWP100) from the IPCC Fourth Assessment Report (AR4) (United Nations Climate Change 2026).[52]

B.3.1 Indirect Emission Intensities

Leontief inverse matrix. Let $\mathbf{I} \in \mathbb{R}^{K \cdot P \times K \cdot P}$ denote the identity matrix, $\mathbf{A} \in \mathbb{R}^{K \cdot P \times K \cdot P}$ the matrix of technical input coefficients from EXIOBASE, $\mathbf{x} \in \mathbb{R}^{K \cdot P \times 1}$ total economic output, and $\mathbf{y} \in \mathbb{R}^{K \cdot P \times 1}$ final demand. The standard input-output system is

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{y}.$$

Assuming $(\mathbf{I} - \mathbf{A})$ is invertible, this implies

$$\mathbf{x} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{y} \equiv \mathbf{L}\mathbf{y},$$

where $\mathbf{L} \equiv (\mathbf{I} - \mathbf{A})^{-1}$ is the Leontief inverse. The Leontief inverse describes the total input required from each sector-region combination to produce one unit of final demand in another sector-region combination.

[51] Measuring emissions in (CO_2e) rather than in physical units of (CO_2) does not affect the essence of our quantitative results as CO_2 accounts for 88% of total greenhouse gas emissions in Germany (Umweltbundesamt 2025b, German Statistical Office 2025c) and the share relative to other greenhouse gases is approximately constant along the income distribution.

[52] HFC and PFC emissions are aggregated in EXIOBASE v3.9.6 and reported directly in kg CO_2e.

Emission coefficients. We derive product-region emission coefficients from EXIOBASE's satellite accounts. Let $\mathbf{F} \in \mathbb{R}^{20 \times K \cdot P}$ collect total emissions of the 20 greenhouse gases falling under the Kyoto protocol by sector-region combination and let $\mathbf{g} \in \mathbb{R}^{1 \times 20}$ denote the vector of GWP100 characterization factors, where entry g_i converts emissions of greenhouse gas i into CO_2e equivalents. The vector of emission coefficients is given by

$$\mathbf{E} = \mathbf{g} \cdot \left(\mathbf{F} \cdot \mathrm{diag}(\mathbf{x})^{-1}\right) \quad \in \mathbb{R}^{1 \times K \cdot P},$$

where $\mathrm{diag}(\mathbf{x})$ denotes the diagonal matrix with the entries of $\mathbf{x}$ on its main diagonal. The entry E_{kp} gives the total CO_2e emissions per unit of output for sector-region pair (k, p).

We set $E_r = 0$ for all sector-region combinations r with $x_r <$ 1e6 to avoid division by zero or unrealistically high emission intensities arising from very low levels of economic output. We interpret these cases as likely driven by measurement error rather than genuine production relationships (Hardadi, Buchholz, and Pauliuk 2021).

Matrix of regional demand shares. Let $\mathbf{Y^{HH}} = [y_{kp}^{HH}] \in \mathbb{R}^{K \times P}$ denote the matrix of final household demand in Germany, where y_{kp}^{HH} is household final demand for product p in region k. For each product p, define the regional demand shares as

$$\omega_{kp} \equiv \begin{cases} \frac{y_{kp}^{HH}}{\sum_{k'=1}^{K} y_{k'p}^{HH}} & \text{if } \sum_{k'=1}^{K} y_{k'p}^{HH} > 0 \\ 0 & \text{else} \end{cases} .$$

Intuitively, ω_{kp} gives region k's share of Germany's household final demand for product p. Let $\mathbf{\Omega} = [\omega_{kp}] \in \mathbb{R}^{K \times P}$ denote the resulting matrix of regional demand shares. We then transform $\mathbf{\Omega}$ into a stacked diagonal-block matrix

$$\tilde{\boldsymbol{Y}} \equiv \mathbf{\Omega} * \mathbf{I} = \begin{bmatrix} \mathrm{diag}(\mathbf{\Omega}_{1,:}) \\ \mathrm{diag}(\mathbf{\Omega}_{2,:}) \\ \vdots \\ \mathrm{diag}(\mathbf{\Omega}_{K,:}) \end{bmatrix} \quad \in \mathbb{R}^{KP \times P},$$

where $\mathbf{I} \in \mathbb{R}^{P \times P}$ is the identity matrix and $*$ denotes the Khatri-Rao product. The indirect emission intensity vector at the basic-price product level is then given by $\mathbf{f}_{\text{ind}}^{\text{ISIC}} = \mathbf{EL\tilde{Y}}$.

B.3.2 Conversion from Basic to Purchaser Prices

Household consumption expenditures in the EVS are reported at purchaser prices, whereas $\mathbf{f}_{\mathbf{ind}}^{\mathbf{ISIC}}$ is defined at basic prices. To align the two concepts, we convert ISIC-level indirect emission intensities from basic to purchaser prices using the trade margins and net taxes on products reconstructed from the German supply-use and input-output tables and mapped into the EXIOBASE framework.[53]

For each product p, one euro of expenditure at purchaser prices can be decomposed into net product taxes, trade margins, and the basic price value of the product:

$$1 = \tau_p + (1-\tau_p)\mu_p + (1-\tau_p)(1-\mu_p),$$

where τ_p denotes the product tax rate (net of subsidies) and μ_p the trade margin ratio.[54] Thus, τ_p is the share of one euro of expenditure at purchaser prices that goes to taxes, $(1-\tau_p)\mu_p$ is the share paying for wholesale and retail trade services, and $(1-\tau_p)(1-\mu_p)$ is the share accruing to the producer at basic prices. Since taxes are a pure price wedge, they do not generate emissions. The remaining two components instead inherit the indirect emission intensities of trade services and the production of the product p, respectively.

Let $\mathcal{S}$ denote the set of EXIOBASE products corresponding to wholesale and retail trade services.[55] For products outside this set, $p \notin \mathcal{S}$, the indirect emission intensity per euro of purchaser price expenditure is given by

$$\tilde{f}^{\text{ISIC}}_{\text{ind},p} = (1-\tau_p)\left[(1-\mu_p)f^{\text{ISIC}}_{\text{ind},p} + \mu_p \bar{f}_{\text{trade}}\right] \quad \text{if } p \notin \mathcal{S},$$

where

$$\bar{f}_{\text{trade}} = \sum_{s\in\mathcal{S}} \gamma_s f^{\text{ISIC}}_{\text{ind},s} \qquad \text{with} \quad \gamma_s = \frac{|\hat{m}_s|}{\sum_{s'\in\mathcal{S}} |\hat{m}_{s'}|},$$

denotes the trade margin weighted average indirect emission intensity of wholesale and retail trade services, and m_s denotes the trade margin attributed to trade service product s[56]. The intuition is straightforward: out of one euro spent on product p at purchaser prices, the share τ_p is absorbed by product taxes and is therefore emission neutral. Of the remaining post-tax expenditure, the fraction $(1-\mu_p)$ accrues to the producer of product p at basic prices, implying that the share $(1-\tau_p)(1-\mu_p)$ of the original purchaser price euro inherits the emission intensity $f^{ISIC}_{\text{ind},p}$. The complementary share $(1-\tau_p)\mu_p$ pays for the wholesale and retail trade

[53] The most recent version of EXIOBASE does not contain supply-use tables and therefore does not provide the valuation information required for this conversion. We reconstruct the corresponding trade margin and net product taxes using the German input-output model of the German Federal Statistical Office. See Appendix B.2.2 for details.

[54] We construct both the net product tax rate τ_p and the trade margin ratio μ_p from the EXIOBASE use tables. Specifically, we compute $\tau_p = \hat{t}^{ISIC,200}_p / y^{ISIC,200}_{pp,p}$ and $\mu_p = \hat{m}^{ISIC,200}_p / (y^{ISIC,200}_{pp,p} - \hat{t}^{ISIC,200}_p)$, where $\hat{t}^{ISIC,200}_p$ denotes net product taxes, $\hat{m}^{ISIC,200}_p$ trade margins, and $y^{ISIC,200}_{pp,p}$ household final demand at purchaser prices, respectively. We set $\tau_p = 0$ and $\mu_p = 0$ whenever the corresponding denominator is zero.

[55] Wholesale and retail trade service products are defined as EXIOBASE products with negative balanced trade margins, i.e. $\hat{m}^{EXIO,200}_p < 0$. This yields five products: "Retail trade services of motor fuel", "Wholesale trade and commission trade services", "Sale, maintenance & repair of motor vehicles, parts & motorcycles", "Retail trade services; repair of personal & household goods" and "Distribution services of gaseous fuels through mains". The last three products play a dual role, as they provide trade services while also exhibiting positive household final demand.

[56] See Section B.2.2 for details on the construction of the trade margin in EXIOBASE.

services embodied in the trade margin and therefore inherits the average trade service emission intensity $\tilde{f}^{\text{ISIC}}_{\text{trade}}$. Hence, $\tilde{f}^{\text{ISIC}}_{\text{ind},p}$ is a weighted average of the indirect emission intensities embodied in the production of product p and in the associated trade services, with weights given by their respective shares in one euro of purchaser-price expenditure.

For EXIOBASE products corresponding to wholesale and retail trade services $p \in \mathcal{S}$, only the part of household demand that reflects direct final consumption should be assigned to the product's own indirect emission intensity. Define the intermediation share of trade-service product p as

$$\alpha_p = \frac{|\hat{m}_p|}{\sum_{k'=1}^{K} y^{HH}_{k'p}}.$$

Intuitively, α_p measures the share of household final demand for product p that is already embodied in the trade margins allocated to other products. The purchaser price indirect emission intensity of product $p \in \mathcal{S}$ is then given by

$$\tilde{f}^{\text{ISIC}}_{\text{ind},p} = (1-\alpha_p)(1-\tau_p)(1-\mu_p) f^{\text{ISIC}}_{\text{ind},p} \quad \text{if } p \in \mathcal{S},$$

The factor $1 - \alpha_p$ captures the part of product p's output that reflects direct household final demand rather than trade intermediation services already embodied in the trade margins assigned to other products. This distinction is relevant for trade service sectors that are not pure intermediaries. If a trade service sector is used purely as an intermediary, then $\alpha_p = 1$, and its purchaser price indirect emission intensity is zero, since all of its supply-chain emissions are already attributed through the trade margins of the producing products.

B.3.3 Mapping from EXIOBASE to COICOP Classification

Household consumption expenditures in the EVS are recorded in the COICOP classification, whereas EXIOBASE is organized according to the ISIC classification. To combine the two data sources, we map indirect emission intensities from the EXIOBASE to the COICOP classification using a correspondence matrix. Let $\mathbf{M} \in \mathbb{R}^{P \times N}$ denote the correspondence matrix, where the n-th column allocates one euro of purchaser price expenditure in COICOP category n across EXIOBASE products. The resulting vector of indirect emission intensities in the COICOP classification is then given by

$$\tilde{\mathbf{f}}^{\text{COICOP}}_{\text{ind}} = \tilde{\mathbf{f}}^{\text{ISIC}}_{\text{ind}} \cdot \mathbf{M} \quad \in \mathbb{R}^{1 \times N},$$

where $\tilde{f}^{\text{COICOP}}_{\text{ind},n}$ measures the indirect CO_2e emissions embodied in one euro of purchaser price expenditure in COICOP category n.

We construct the correspondence matrix $\mathbf{M}$ in three steps. First, we define an initial binary correspondence between EXIOBASE and COICOP products. Second, we apply the RAS method to adjust this matrix so that its row and column totals match household final demand in each EXIOBASE and COICOP product, respectively. Third, we normalize the resulting matrix column-wise so that each column sums to one. The final step converts the adjusted expenditure flows into purchaser price expenditure weights, where each entry M_{pn} gives the weight placed on EXIOBASE product p when constructing the indirect emission intensity of COICOP category n. The resulting matrix therefore defines a many-to-many mapping, whereby each COICOP category can be associated with multiple EXIOBASE products.

Initial correspondence matrix. Let $\mathbf{A}_0 = [a_{ij}] \in \mathbb{R}^{P \times N}$ denote the initial binary correspondence matrix, where $a_{ij} = 1$ if COICOP product j is assigned to EXIOBASE product i (based on product similarities), and 0 otherwise. We construct $\mathbf{A}_0$ from the correspondence proposed by Stadler et al. (2018) and Hardadi, Buchholz, and Pauliuk (2021) with additional modifications to ensure convergence of the RAS algorithm.[57]

RAS method. Let $\mathbf{r} \in \mathbb{R}^{P \times 1}$ denote household final demand by EXIOBASE product, expressed at purchaser prices, and let $\mathbf{s} \in \mathbb{R}^{N \times 1}$ denote aggregate household expenditure by COICOP category in the EVS.[58] The RAS method adjusts $\mathbf{A}_0$ iteratively as follows:

1. **Row Adjustment**: The rows of the correspondence matrix $\mathbf{A}^{k-1}$ at iteration k are scaled to match the final demand in each EXIOBASE product, given by the column vector $\mathbf{r} \in \mathbb{R}^{P \times 1}$. The updated correspondence matrix is given by

$$\mathbf{A}^k = \hat{\mathbf{r}}^k \mathbf{A}^{k-1},$$

where the diagonal matrix of row scaling factors $\hat{\mathbf{r}}^k \in \mathbb{R}^{P \times P}$ is defined as

$$\hat{\mathbf{r}}^k = \operatorname{diag}\left(\left[\frac{r_i}{\sum_{j=1}^{N} \mathbf{A}_{ij}^{k-1}} \right]_{i=1}^{P} \right)$$

with r_i being the i-th element of vector $\mathbf{r}$ and $\sum_{j=1}^{N} \mathbf{A}_{ij}^{k-1}$ being the row sum of matrix $\mathbf{A}^{k-1}$.[59] Note that at the first iteration ($k = 1$), the matrix $\mathbf{A}^{k-1}$ is given by the initial correspondence matrix A_0.

[57] Specifically, we augmented the support of A_0 as follows. We manually added the links required for block feasibility, meaning that row and column totals balance within each connected support block. We then added further links identified by a maximum-flow feasibility test, which checks whether EXIOBASE household final demand for each ISIC product can be allocated across the admissible COICOP categories while simultaneously matching EVS expenditure in every COICOP category. Each of these added links was reviewed case by case for economic plausibility.

[58] Because EXIOBASE reports final demand at basic prices, we first convert it to purchaser prices using EXIOBASE use table. See Appendix B.2.2 for details on the construction of the EXIOBASE use table. We scale aggregate EXIOBASE household final demand to match the corresponding national accounts aggregate, so that the row and column targets are mutually consistent.

[59] Note that this requires $\sum_{j=1}^{N} \mathbf{A}_{ij}^{k-1} \neq 0 \ \forall \ i$.

2. **Column Adjustment**: The columns of the correspondence matrix $\mathbf{A}^{k-1}$ at iteration k are scaled to match the total expenditures in each COICOP category from EVS, given by the column vector $\mathbf{s} \in \mathbb{R}^{N\times 1}$. The updated correspondence matrix is given by

$$\mathbf{A}^{k+1} = \mathbf{A}^k \hat{\mathbf{s}}^k = \hat{\mathbf{r}}^k \mathbf{A}^{k-1} \hat{\mathbf{s}}^k,$$

where the diagonal matrix of column scaling factors $\hat{\mathbf{s}}^k \in \mathbb{R}^{N\times N}$ is defined as

$$\hat{\mathbf{s}}^k = \operatorname{diag}\left(\left[\frac{s_j}{\sum_{i=1}^{P} \mathbf{A}_{ij}^k}\right]_{j=1}^{N}\right)$$

with s_j being the j-th element of vector $\mathbf{s}$ and $\sum_{i=1}^{P} \mathbf{A}_{ij}^k$ being the column sum of the matrix $\mathbf{A}^k$.[60]

These two steps are repeated iteratively until convergence. Specifically, convergence is achieved when

$$\max\left\{ \max_{i=1,\ldots,P} \left| \frac{\sum_{j=1}^{N} \mathbf{A}_{ij}^k - r_i}{r_i} \right|_{i=1}^{P}, \max_{j=1,\ldots,N} \left| \frac{\sum_{i=1}^{P} \mathbf{A}_{ij}^k - s_j}{s_j} \right|_{j=1}^{N} \right\} < \epsilon_{tol},$$

where $\epsilon_{tol} = 10^{-6}$. Rows and columns with zero target values ($r_i = 0$ or $s_j = 0$) are excluded from the maximum operators, as their corresponding fitted entries deterministically scale to zero. This ensures that both the row and column sums of $\mathbf{A}^{k+1}$ match r and s. Once the algorithm converges, we define the resulting correspondence matrix as $\tilde{\mathbf{A}} \in \mathbf{R}^{P\times N}$.

Normalization of correspondence matrix. Finally, we normalize each column of $\tilde{\mathbf{A}}$ to sum to one:

$$\mathbf{M} = \tilde{\mathbf{A}} \oslash (\mathbf{1}^{\mathsf{T}} \tilde{\mathbf{A}}),$$

where $\oslash$ denotes element-wise division and $\mathbf{1} \in \mathbb{R}^{P\times 1}$ is a column vector of ones. Conceptually M_{pn} gives the share of one euro spent in COICOP category n that is allocated to EXIOBASE product p. The matrix $\mathbf{M}$ therefore defines a many-to-many mapping from COICOP categories to EXIOBASE products.

B.3.4 Direct Emission Intensities

We compute direct carbon emission intensities for each COICOP product by converting physical emission factors for the underlying fuels into monetary emission intensities. Direct emissions are assigned only to COICOP categories that generate emissions at the point of household use: gas (c04.5.2), liquid fuels (c04.5.3), solid fuels (c04.5.4), and fuels and lubricants for personal transport equipment (c07.2.2). For each COICOP product n, let K_n denote the set of underlying

[60] Note that this requires that $\sum_{i=1}^{P} \mathbf{A}_{ij}^k \neq 0 \ \forall\, j$.

fuels and let $k \in K_n$ index individual fuels. We then compute the monetary direct emission intensity for each fuel k as

$$\tilde{f}_{dir,k} = \frac{f^{physical}_{dir,k}}{p^{2018}_k},$$

where $f^{physical}_{dir,k}$ is the direct emission factor per physical unit consumed and p^{2018}_k is the corresponding 2018 consumer price per physical unit. This conversion yields direct emissions per euro of expenditure. Table 5 reports, for each fuel within the relevant COICOP categories, the physical emission factor measured in tCO_2e/TJ and the corresponding monetary emission intensity.[61]

Finally, we aggregate fuel specific monetary emission intensities to the COICOP level by taking a weighted average across the fuels $k \in K_n$ underlying COICOP product n:

$$\tilde{f}^{COICOP}_{dir,n} = \sum_{k \in K_n} s_{k,n} \tilde{f}^{dir}_k,$$

where $s_{k,n}$ denotes the share of fuel k in household fuel consumption within COICOP product n. For gas and solid fuels, we construct these shares using the 2011–2013 German Residential Energy Consumption Survey (GRECS), provided by RWI – Leibniz-Institut für Wirtschaftsforschung.[62] For fuels and lubricants for personal transport equipment, we use the corresponding fuel shares from the German Statistical Office (Brockjan, Maier, Kott, and Sewald 2021). For COICOP categories that contain only one relevant fuel type, we set $s_{k,n} = 1$. Table 5 reports the resulting consumption shares. The resulting vector of direct emission intensities is denoted by $\tilde{\mathbf{f}}^{COICOP}_{dir} \in \mathbb{R}^{N \times 1}$. For all COICOP categories that do not generate direct emissions at the household level, the corresponding entries of $\tilde{\mathbf{f}}^{COICOP}_{dir}$ are set to zero.

B.4 Validation of Emission Intensities

Table 6 reports the implied aggregate emissions of German households and compares them with the estimates reported by German Statistical Office (2025b). Since the benchmark statistics are reported for CO_2, while our main estimates are expressed in CO_2 equivalents, the table reports both CO_2e and CO_2 estimates. Our estimates imply total annual household emissions of 731.3 $MtCO_2e$, of which 32% are attributable to direct emissions. Dividing aggregate emissions by the number of tax units gives an average carbon footprint of 16.9 tCO_2e per tax unit. Our

[61] For each fuel, we collect greenhouse gas specific physical emission factors for CO_2, CH_4, and N_2O. We then convert CH_4 and N_2O into CO_2-equivalents using the corresponding characterization factors and sum across gases to obtain the total physical emission factor in tCO_2e/TJ. We exclude SF_6, HFCs, and PFCs, as these greenhouse gases do not arise as direct household emissions for the fuels considered.

[62] The GRECS data is a representative panel on the energy consumption of German households (RWI and forsa 2015). We use the 2013 wave to construct consumption-weighted fuel shares based on standardized household fuel consumption measured in kWh and the corresponding fuel-specific survey weights.

Energy Source	**Physical Emission Factor** (tCO_2e/TJ)	**Emission Intensity** ($kgCO_2e/EUR$)	**Share**
Gas (c04.5.2)		**3.00**	
Liquefied Petroleum Gas (LPG)	66.45	3.37	0.045
Natural Gas	55.82	2.98	0.955
Liquid Fuels (c04.5.3)		**3.87**	
Heating Oil	74.18	3.87	1.000
Solid Fuels (c04.5.4)		**8.00**	
Hard Coal	102.58	5.57	0.010
Lignite	104.59	7.95	0.016
Wood	105.60	8.14	0.852
Wood Pellets	101.14	7.22	0.122
Fuels for Personal Transport Equipment (c07.2.2)		**1.85**	
Diesel Fuel	76.04	2.09	0.422
Gasoline	75.09	1.68	0.578

Table 5: Physical and Monetary Direct Emission Intensities

Notes: Physical emission factors for CO_2 are taken from Umweltbundesamt (2025a). Physical emission factors for CH_4 and N_2O are taken from Blank, Wickert, Obermeier, and Friedrich (1999), except for liquefied petroleum gas (LPG), for which we use IPCC (2006a), and for diesel fuel and gasoline, for which we use IPCC (2006b). We convert CH_4 and N_2O into CO_2-equivalents using the corresponding characterization factors and sum across gases to obtain the total physical emission factor in tCO_2e/TJ. Fuel shares for gas (c04.5.2) and solid fuels (c04.5.4) are based on the 2013 wave of the 2011–2013 German Residential Energy Consumption Survey (GRECS) (RWI and forsa 2015), provided by RWI – Leibniz-Institut für Wirtschaftsforschung, and are computed as consumption-weighted shares using standardized household fuel consumption measured in kWh and the corresponding fuel-specific survey weights. Fuel shares for fuels and lubricants for personal transport equipment (c07.2.2) are taken from the German Statistical Office (Brockjan, Maier, Kott, and Sewald 2021). For liquid fuels (c04.5.3), the fuel share is set to one.

estimates in CO_2 are of the same order of magnitude as the estimates reported by the German statistical office.

	Own Estimates		German Statistical Office
	CO_2e	CO_2	CO_2
Total Emissions	**731.26 Mt**	**603.57 Mt**	**592.70 Mt**
Indirect Emissions	500.12 Mt	375.33 Mt	376.72 Mt
Direct Emissions	231.14 Mt	228.23 Mt	215.97 Mt
Carbon Footprint per Tax Unit	**16.92 t**	**13.96 t**	**13.71 t**

Table 6: Aggregate GHG emissions of German households.

Notes: Own estimates are obtained by multiplying COICOP-level emission intensities by aggregate household consumption expenditures in purchaser prices obtained from the EVS and summing across COICOP categories. The estimates of the German Statistical office are obtained from German Statistical Office (2025b). Indirect emissions refer to emissions generated along the production chain, while direct emissions are released directly by households during consumption. Per tax unit values are obtained by dividing aggregate emissions by the number of tax units.

C Micro Data

We use two primary data sources: (i) detailed income and consumption expenditure data from the *Einkommens- und Verbrauchsstichprobe* (EVS), provided by the German Statistical

Office (Research Data Centre of the Federal Statistical Office and the statistical offices of the Länder (RDC) 2018b, Research Data Centre of the Federal Statistical Office and the statistical offices of the Länder (RDC) 2018a) and (ii) German administrative income tax data from the *Lohn- und Einkommensteuerstatistik* (LESt), provided by the German Federal Statistical Office (Research Data Centre of the Federal Statistical Office and the statistical offices of the Länder (RDC) 2018c). Both datasets refer to the year of 2018. We limit our sample to the working-age population, defined as tax units in which the main earner is between 20 and 60 years old.

EVS. The EVS provides quarterly household-level data on consumption expenditures, gross labor income, and socioeconomic characteristics. We convert all monetary variables to an annual level by multiplying the quarterly values by four. The unit of observation is the household. To ensure comparability with income tax data from the LESt, we convert the data from the household level to the tax unit level. Specifically, we split non-married couples into separate tax units by allocating household consumption expenditures according to each person's share of household labor income. This leaves us with a sample of around 26 000 tax units.

Consumption expenditures are reported in purchaser's prices and classified according to the five-digit COICOP system. For tractability, we aggregate expenditures to the four-digit COICOP level by summing across all corresponding five-digit subcategories, resulting in 107 different products and services.[63] Table 7 documents the aggregation and provides the relevant variables in the EVS.

We define gross labor income as the sum of income from employment and self-employment. Accordingly, we assign each tax unit to the corresponding percentile of the gross labor income distribution obtained from the LESt. Transfer payments include unemployment benefits (*Arbeitslosengeld II; Sozialgeld nach Sozialgesetzbuch II*), employment promotion transfers, housing allowance, social assistance and child supplement.[64]

To address the well-known issue of underreporting in household survey data, we scale consumption expenditures in the EVS by a constant factor such that the population-weighted aggregate consumption in the EVS aligns with national accounts data from German Statistical Office (2024a). Similarly, we adjust transfer receipts by applying transfer specific scaling factors to ensure consistency with administrative aggregates. Specifically, unemployment benefits are scaled to match total reported spending from Bundesagentur für Arbeit (2018), housing

[63] Detailed expenditure data for food (c01.1), non-alcoholic (c01.2) and alcoholic beverages (c02.1) is not available in the main EVS file (gf5). Therefore, we rely on supplementary data from the gf4 file, which contains a detailed recording booklet for food, beverages and tobacco products, covering roughly 20% of the tax units included in the main EVS file (gf5). We compute the expenditure share of each product within the COICOP groups food (c01.1), non-alcoholic beverages (c01.2) and alcoholic beverages (c02.1) at the tax unit level. These shares are then averaged across tax units within net income bins, using survey weights. Finally, we impute consumption expenditures in the gf5 file by multiplying total expenditure in each COICOP group (EVS variables EF242, EF243, and EF244) by the corresponding average expenditure share.

[64] Gross labor income is computed as the sum of the following EVS variables: EF109U - EF115U, EF118U - EF137U, and EF176. Transfer payments are computed as the sum of the following EVS variables: EF149U, EF150U1, EF151U, EF152U, EF154U - EF157U, and EF168U.

allowances to aggregates from German Statistical Office (2024b), social assistance to aggregates from German Statistical Office (2025a), and child supplements to aggregates from Familienkasse Direktion SR1 (2018).[65]

LESt. The LESt comprises a 10% cross-sectional random sample of German taxpayers based on administrative income tax records. We use data on annual gross labor income to calibrate the income distribution. Note that the unit of observation is the tax unit. This implies that if a married couple files taxes separately, we cannot identify them as belonging to the same household in our data. In 2018, around 8.1% of married couples filed taxes separately. Labor income of a tax unit is defined as the sum of income from employment and self-employment. Tax payments are calculated as the sum of the assessed income tax (*festzusetzende Einkommensteuer*), the solidarity surcharge (*Solidaritätszuschlag*) and the church tax.[66] Our sample includes around 2.28 million tax units.

Code	Product	EVS Variables	Brown/Green	Emission Intensity
		Food		
c01.1.1	Cereals and cereal products	EF58U2 - EF87U2	Green	0.3927 kgCO2e/€
c01.1.2	Live animals, meat and other parts of slaughtered land animals	EF88U2 - EF111U2	Brown	0.6383 kgCO2e/€
c01.1.3	Fish and other seafood	EF112U2 - EF125U2	Brown	0.5273 kgCO2e/€
c01.1.4	Milk, other diary products and eggs	EF126U2 - EF144U2	Brown	0.6412 kgCO2e/€
c01.1.5	Oils and fats	EF145U2 - EF153U2	Brown	0.6366 kgCO2e/€
c01.1.6	Fruits and nuts	EF154U2 - EF177U2	Green	0.4076 kgCO2e/€
c01.1.7	Vegetables, tubers, plantains, cooking bananas and pulses	EF178U2 - EF208U2	Green	0.4078 kgCO2e/€
c01.1.8	Sugar, confectionery and desserts	EF209U2 - EF222U2	Brown	0.6271 kgCO2e/€
c01.1.9	Ready-made food and other food products	EF224U2 - EF249U2	Green	0.4014 kgCO2e/€
c01.2.1	Coffee, tea and cocoa	EF250U2 - EF255U2	Green	0.4210 kgCO2e/€
c01.2.2	Mineral waters, soft drinks, fruit and vegetable juices	EF256U2 - EF267U2	Green	0.3714 kgCO2e/€
c02.1.1	Spirits	EF268U2 - EF271U2	Green	0.3767 kgCO2e/€
c02.1.2	Wine	EF272U2 - EF283U2	Green	0.3743 kgCO2e/€
c02.1.3	Beer	EF284U2 - EF288U2	Green	0.3853 kgCO2e/€
		Clothing		
c03.1.1	Clothing materials	EF247	Brown	0.3240 kgCO2e/€
c03.1.2	Garments	EF248 - EF250	Brown	0.2857 kgCO2e/€
c03.1.3	Other articles of clothing and clothing accessories	EF251	Brown	0.2914 kgCO2e/€
c03.1.4	Cleaning, repair and hire of clothing	EF252, EF253	Green	0.1199 kgCO2e/€
c03.2.1	Shoes and other footwear	EF254 - EF257	Brown	0.2530 kgCO2e/€
c03.2.2	Repair and hire of footwear	EF258	Green	0.1199 kgCO2e/€
c12.3.1	Jewelry, clocks and watches	EF443, EF444	Brown	0.2557 kgCO2e/€
c12.3.2	Other personal effects	EF445	Brown	0.3083 kgCO2e/€
		Housing		
c04.1.1	Actual rentals paid by tenants	EF259 - EF262	Green	0.0751 kgCO2e/€
c04.1.2	Other actual rentals	EF263 - EF267	Green	0.0462 kgCO2e/€
c04.2.1	Imputed rentals of owner-occupiers	EF268 - EF271	Green	0.0450 kgCO2e/€
c04.2.2	Other imputed rentals	EF272 - EF276	Green	0.0462 kgCO2e/€
c04.3.1	Materials for the maintenance and repair of the dwelling	EF277 - EF280	Brown	0.4301 kgCO2e/€
c04.4.1	Water supply	EF285 - EF287	Brown	0.3581 kgCO2e/€
c04.4.2	Refuse collection	EF288 - EF290	Brown	0.4023 kgCO2e/€
c04.4.3	Sewage collection	EF291 - EF293	Brown	0.4023 kgCO2e/€
c04.4.5	Cold operating costs, additional costs	EF297 - EF307	Brown	0.3465 kgCO2e/€
		Electricity		
c04.5.1	Electricity	EF313	-	1.7913 kgCO2e/€
		Heating		
c04.5.2	Gas	EF314 - EF317	Green	3.4322 kgCO2e/€
c04.5.3	Liquid fuels	EF318 - EF320	Green	4.4227 kgCO2e/€

[65] We do not scale employment promotion transfer due to lack of aggregate data.

[66] Gross labor income is calculated as the sum of the following LESt variables: C65120, C65100, C65140, C65163 and C65164. Income tax payments are computed as the sum of the variables C65613, C66104 and C66105. The solidarity surcharge is computed based on the assessed income tax, taking into account the exemption thresholds applicable in 2018.

c04.5.4	Solid fuels	EF321	Brown	8.6185 kgCO2e/€
c04.5.5	Heat energy	EF322 - EF324	Green	4.5929 kgCO2e/€
	Household and Personal Care			
c05.1.1	Furniture and furnishing	EF326, EF327	Brown	0.2890 kgCO2e/€
c05.1.2	Carpets and other floor coverings	EF328, EF329	Brown	0.2994 kgCO2e/€
c05.1.3	Repair of furniture, furnishings and floor coverings	EF330	Green	0.1010 kgCO2e/€
c05.2.0	Household textiles	EF331, EF332	Brown	0.3470 kgCO2e/€
c05.3.1	Major household appliances whether electric or not	EF333 - EF336	Green	0.1741 kgCO2e/€
c05.3.2	Small electric household appliances	EF337	Brown	0.3188 kgCO2e/€
c05.3.3	Repair of household appliances	EF338	Brown	0.2692 kgCO2e/€
c05.4.0	Glassware, tableware and household utensils	EF339, EF340	Brown	0.2791 kgCO2e/€
c05.5.1	Major tools and equipment	EF341, EF342	Green	0.1126 kgCO2e/€
c05.5.2	Small tools and miscellaneous accessories	EF343 - EF345	Brown	0.4548 kgCO2e/€
c05.6.1	Non-durable household goods	EF346	Brown	0.2872 kgCO2e/€
c06.1.1	Pharmaceutical products	EF349 - EF353	Brown	0.2765 kgCO2e/€
c06.1.2	Other medical products	EF354 - EF358	Brown	0.2860 kgCO2e/€
c06.1.3	Therapeutic appliances and equipment	EF359 - EF362	Green	0.1384 kgCO2e/€
c09.5.4	Stationery and drawing materials	EF422	Brown	0.2916 kgCO2e/€
c12.1.2	Electric appliances for personal care	EF438	Green	0.1453 kgCO2e/€
c12.1.3	Other appliances, articles and products for personal care	EF439 - EF441	Brown	0.2905 kgCO2e/€
	Services			
c04.3.2	Services for the maintenance and repair of the dwelling	EF281 - EF284	Brown	0.3295 kgCO2e/€
c04.4.4	Other services relating to the dwelling	EF294 - EF296	Brown	0.3418 kgCO2e/€
c05.6.2	Domestic services and household services	EF347, EF348	Brown	0.2182 kgCO2e/€
c06.2.1	Medical services	EF363	Green	0.1048 kgCO2e/€
c06.2.2	Dental services	EF364	Green	0.1048 kgCO2e/€
c06.2.3	Paramedical services	EF365, EF366	Green	0.1241 kgCO2e/€
c06.3.0	Hospital services	EF367	Green	0.1267 kgCO2e/€
c07.2.4	Other services in respect of personal transport equipment	EF378	Green	0.1557 kgCO2e/€
c08.1.0	Postal services	EF385	Green	0.0866 kgCO2e/€
c08.3.0	Telephone and telefax services	EF387 - EF392	Green	0.0890 kgCO2e/€
c09.1.5	Repair of audio-visual, photographic and information processing equipment	EF398	Brown	0.3380 kgCO2e/€
c10.1.0	Pre-primary and primary education	EF425 - EF427	Green	0.0661 kgCO2e/€
c10.2.0	Secondary education	EF428	Green	0.0661 kgCO2e/€
c10.5.0	Education not definable by level	EF429, EF430	Green	0.0661 kgCO2e/€
c12.1.1	Hairdressing salons and personal grooming establishments	EF434 - EF437	Green	0.0807 kgCO2e/€
c12.4.0	Social protection (private)	EF446 - EF449	Green	0.1155 kgCO2e/€
c12.6.2	Other financial services	EF451	Green	0.0999 kgCO2e/€
c12.7.0	Other services	EF452	Green	0.1030 kgCO2e/€
	Transport and Vacation			
c07.1.1	Motor cars	EF368, EF369	Green	0.1965 kgCO2e/€
c07.1.2	Motor cycles	EF370	Green	0.1959 kgCO2e/€
c07.1.3	Bicycles	EF371	Green	0.3053 kgCO2e/€
c07.1.4	Animal drawn vehicles	EF372	Green	0.2466 kgCO2e/€
c07.2.1	Spare parts and accessories for personal transport equipment	EF374, EF375	Green	0.3088 kgCO2e/€
c07.2.2	Fuels and lubricants for personal transport equipment	EF376	Brown	2.1892 kgCO2e/€
c07.2.3	Maintenance and repair of personal transport equipment	EF377	Green	0.1696 kgCO2e/€
c07.3.1	Passenger transport by railway	EF379	Green	0.2455 kgCO2e/€
c07.3.2	Passenger transport by road	EF380	Green	0.2383 kgCO2e/€
c07.3.3	Passenger transport by air	EF381	Brown	0.8358 kgCO2e/€
c07.3.4	Passenger transport by sea and inland waterways	EF382	Brown	1.6285 kgCO2e/€
c07.3.5	Combined passenger transport	EF383	Brown	0.6750 kgCO2e/€
c07.3.6	Other purchased transport services	EF384	Green	0.1809 kgCO2e/€
c09.6.1	Package holidays (Germany)	EF423	Brown	0.6527 kgCO2e/€
c09.6.2	Package holidays (foreign country)	EF424	Brown	0.8357 kgCO2e/€
c11.2.0	Accommodation services	EF433	Green	0.1914 kgCO2e/€
	Leisure			
c02.2.0	Tobacco	EF245	Green	0.1695 kgCO2e/€
c08.2.0	Telephone and telefax equipment	EF386	Green	0.0896 kgCO2e/€
c09.1.1	Equipment for the reception, recording and reproduction of sound and pictures	EF393, EF394	Brown	0.3037 kgCO2e/€
c09.1.2	Photographic and cinematographic equipment and optical instruments	EF395	Brown	0.3052 kgCO2e/€
c09.1.3	Information processing equipment	EF396	Brown	0.2578 kgCO2e/€
c09.1.4	Recording media	EF397	Brown	0.2979 kgCO2e/€
c09.2.1	Major durables for outdoor recreation	EF399	Brown	0.2758 kgCO2e/€
c09.2.3	Maintenance and repair of other major durables for recreation and culture	EF400	Green	0.1521 kgCO2e/€
c09.3.1	Games, toys and hobbies	EF401	Brown	0.2412 kgCO2e/€
c09.3.2	Equipment for sport, camping and open-air recreation	EF402, EF403	Brown	0.2664 kgCO2e/€
c09.3.3	Gardens, plants and flowers	EF404, EF405	Brown	0.4167 kgCO2e/€
c09.3.4	Pets and related products	EF406	Brown	0.2964 kgCO2e/€
c09.4.1	Recreational and sporting services	EF407 - EF409	Brown	0.1809 kgCO2e/€
c09.4.2	Cultural services	EF410 - EF416	Green	0.0897 kgCO2e/€
c09.4.3	Games of chance	EF417	Green	0.1017 kgCO2e/€
c09.5.1	Books	EF418	Green	0.0986 kgCO2e/€

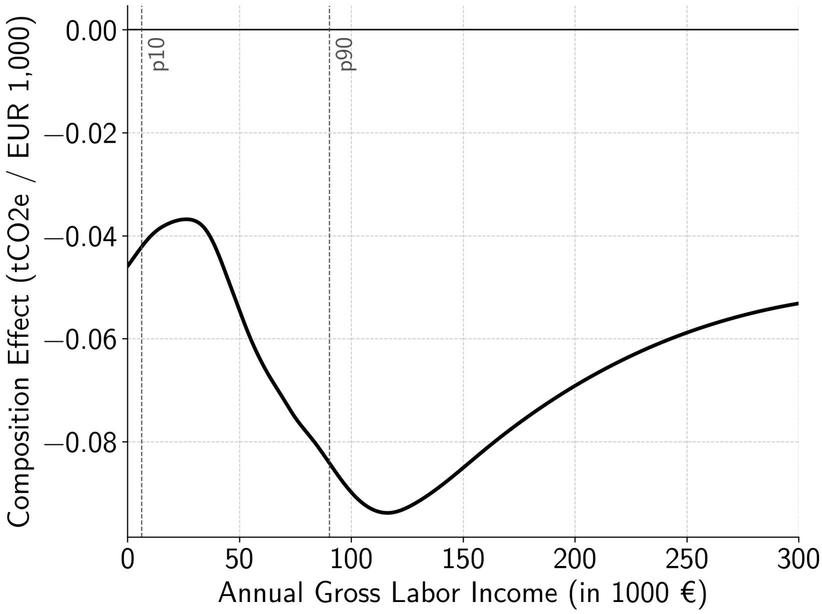

(a) Composition Effect

Figure 9: Composition Component of the Marginal Propensity to Pollute

Notes: The figure shows the composition component of the marginal propensity to pollute, defined as the part of the change in emissions from an additional unit of total expenditure that is due to the induced change in the composition of the consumption basket. It is given by the difference between the MPP, defined as $\mathrm{MPP}(p) = \hat{f}'(e(p))$ and average emission intensity $\iota(p) = \hat{f}(p)/e(p)$, where $\hat{f}(\cdot)$ is a cubic smoothing spline of $\tilde{f}(p)$ on $e(p)$. The difference is computed based on the cubic smoothing spline of both objects. The x-axis shows annual gross labor income of a tax unit. The grey vertical dashed lines mark the 10th and 90th percentiles of the annual gross labor income distribution.

c09.5.2	Newspapers and periodicals	EF419, EF420	Green	0.1173 kgCO2e/€
c09.5.3	Miscellaneous printed matter	EF421	Green	0.0986 kgCO2e/€
c11.1.1	Restaurants, cafés and the like	EF431	Brown	0.2391 kgCO2e/€
c11.1.2	Canteens	EF432	Brown	0.2391 kgCO2e/€

Table 7: Classification of Products

Note: The COICOP groups c01.1 (food), c01.2 (non-alcoholic beverages), and c02.1 (alcoholic beverages) are based on item-level expenditure data from the EVS recording booklet for food, beverages and tobacco products (gf4). All remaining COICOP groups are based on data from the EVS main file (gf5). Emission intensities are expressed in kgCO2e per euro of expenditure. Within each product category, products are classified as either 'brown' or 'green' based on whether their emission intensity is above or below the cutoff emission intensity for that category. The cutoff emission intensity is determined by K-means clustering. The methodology for computing carbon footprints is described in Appendix B.

D Household Carbon Footprint

D.1 Decomposition of Marginal Propensity to Pollute

Figure 9 plots the composition component of the marginal propensity to pollute (MPP) along the income distribution.

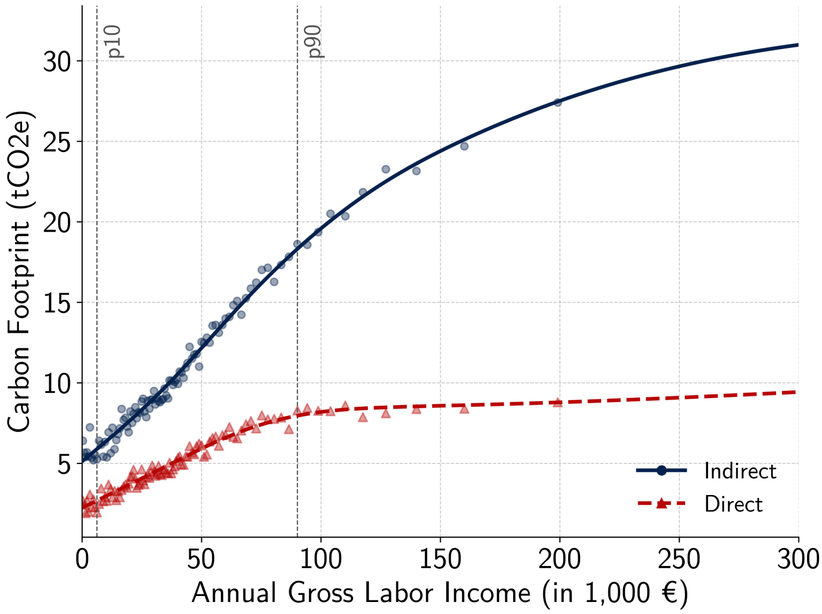

Figure 10: Direct and Indirect Carbon Emissions

Notes: The figure shows the carbon footprint of indirect emissions (blue dots) and direct emissions (red triangles) for each income percentile, both expressed in tCO_2e. The lines show cubic smoothing spline fits to the respective data moments. The x-axis shows annual gross labor income of a tax unit. The grey vertical dashed lines mark the 10th and 90th percentiles of the annual gross labor income distribution.

D.2 Anatomy of the Income Gradient in Carbon Emissions

D.2.1 Direct vs. Indirect Emissions

Figure 10 decomposes the percentile level carbon footprint into indirect and direct emissions. In the aggregate, indirect emissions account for 69% of total household emissions and account for 68% of the gap in carbon emissions between the 10th and 90th percentiles of the income distribution.

D.2.2 Horizontal Carbon Inequality

To assess the role of observable horizontal heterogeneity in carbon footprints, we compute covariate-standardized carbon footprints for each percentile of the gross labor income distribution. We first estimate the following survey-weighted linear projection at the tax-unit level:

$$f_i = \alpha + \sum_{p \neq p_0} \gamma_p \mathbf{1}[p_i = p] + \mathbf{X}_i' \boldsymbol{\beta} + u_i,$$

where f_i denotes the carbon footprint of tax unit i and p_i denotes its income percentile. The vector $\mathbf{X}_i$ contains observed characteristics, including household size, education, homeownership status, heating fuel type, rural residence status and age.

We then compute the covariate-standardized carbon footprint for each income percentile as

$$f_{std}(p) = \hat{\alpha} + \hat{\gamma}_p + \overline{\boldsymbol{X}}'_w \hat{\boldsymbol{\beta}},$$

where $\overline{\boldsymbol{X}}'_w$ is the survey-weighted mean of the covariates in the full estimation sample. Thus, every income percentile is evaluated at the same empirical distribution of observed characteristics.

In contrast, the raw carbon footprint profile is computed as the survey-weighted average carbon footprint among tax units observed in income percentile p and therefore uses the actual covariate distribution of each income percentile.

Figure 11(a) illustrates the raw (blue solid line) and covariate-standardized adjusted carbon footprints (grey dashed line) along the gross labor income distribution. The gap in carbon emissions between the 10th and 90th percentiles falls by about 24%; equivalently, 76% of the raw gap remains. The covariate-standardized profile therefore remains upward sloping, indicating that observable horizontal heterogeneity accounts for part of carbon inequality observed in the data, but a substantial vertical component remains. Figure 11(b) shows the corresponding marginal propensity to pollute, computed from the raw and standardized carbon footprint profiles.

D.2.3 Carbon Footprint by Product Category

Figure 12(a) shows the average carbon footprint by income decile of the annual gross labor income distribution and broad product category. Figure 12(b) reports for each product category i the contribution to the gap in the carbon footprint between the 90th and 10th income percentile, e.g. $\frac{\tilde{f}_i(p90)-\tilde{f}_i(p10)}{\tilde{f}(p90)-\tilde{f}(p10)}$, where $\tilde{f}_i(p)$ is the carbon footprint of product category i of percentile p.

D.3 Aggregation Error

Figure 13 quantifies the aggregation error induced by aggregating the detailed COICOP categories into the model's broader product–variety cells. The left panel compares the carbon footprint calculated at the detailed COICOP level (blue line) with the model counterpart obtained by aggregating COICOP products to broad product–variety cells (grey line), for both the raw and savings-adjusted carbon footprints. Note that the savings adjustment does not affect the aggregation error because the savings adjustment is a proportional scalar applied identically to both the detailed and the aggregate estimate. The right panel shows the resulting aggregation error, defined as the percentage difference between the aggregated and detailed estimates, across income percentiles. The two estimates overlap closely up to around the 95th percentile of the gross labor income distribution. Above the 90th percentile, the aggregated estimate in-

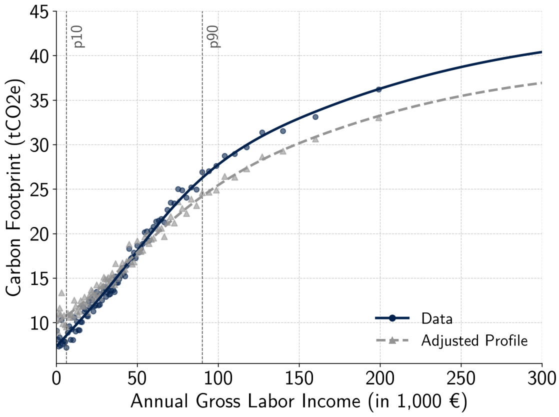

(a) Carbon Footprint

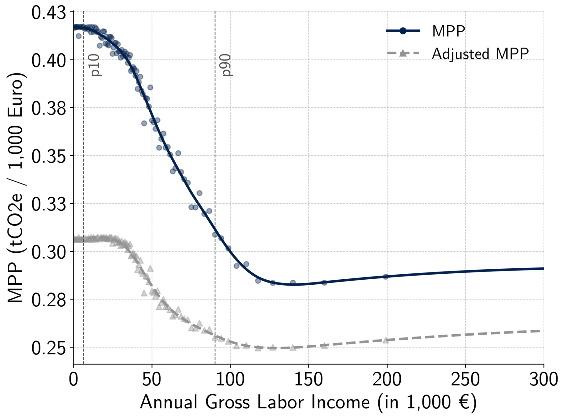

(b) Marginal Propensity to Pollute

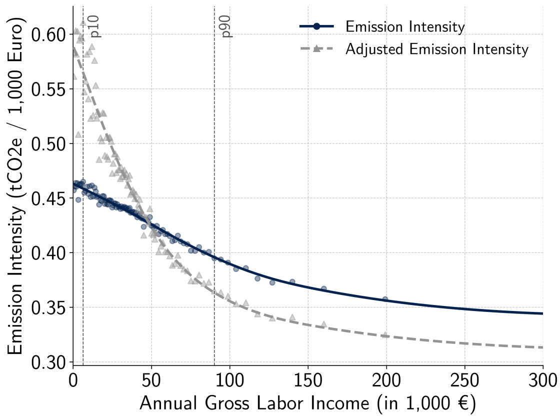

(c) Emission Intensity

Figure 11: Covariate-Standardized Carbon Footprint and Marginal Propensity to Pollute

Notes: The left panel illustrates the carbon footprint $f(p)$ in the data (blue dots) and the covariate-standardized carbon footprint $f_{std}(p)$ (grey triangles) for each income percentile p, both expressed in tCO_2e. The covariate-standardized profile is constructed from a linear projection of tax unit carbon footprints on income percentile fixed effects and observed characteristics, evaluating the controls at their common distribution in the EVS sample for every percentile. The right panel shows the marginal propensity to pollute (MPP) in the data (blue dots) and the covariate-standardized MPP, both expressed in tCO_2e per 1 000 Euro. The MPP is defined as the derivative of the spline-fitted carbon footprint with respect to total expenditure, evaluated at each percentile: $\text{MPP}(p) = \hat{f}'(e(p))$, where $\hat{f}(\cdot)$ is a cubic smoothing spline of $f(p)$ on $e(p)$. The covariate-standardized MPP is defined as $\text{MPP}_{std}(p) = \hat{f}'_{std}(e(p))$, where $\hat{f}_{std}(\cdot)$ is a cubic smoothing spline of $f_{std}(p)$ on $e(p)$. The bottom panel shows the emission intensity $\iota(p) = \hat{f}(p)/e(p)$ (blue dots) and the covariate-adjusted emission intensity is defined as $\iota_{std} = \hat{f}_{std}(p)/e(p)$ (grey dots), both expressed in tCO_2e per 1 000 Euro. In all panels, the lines show cubic smoothing spline fits to the respective data moments. The x-axis shows annual gross labor income of a tax unit. The grey vertical dashed lines mark the 10th and 90th percentiles of the annual gross labor income distribution.

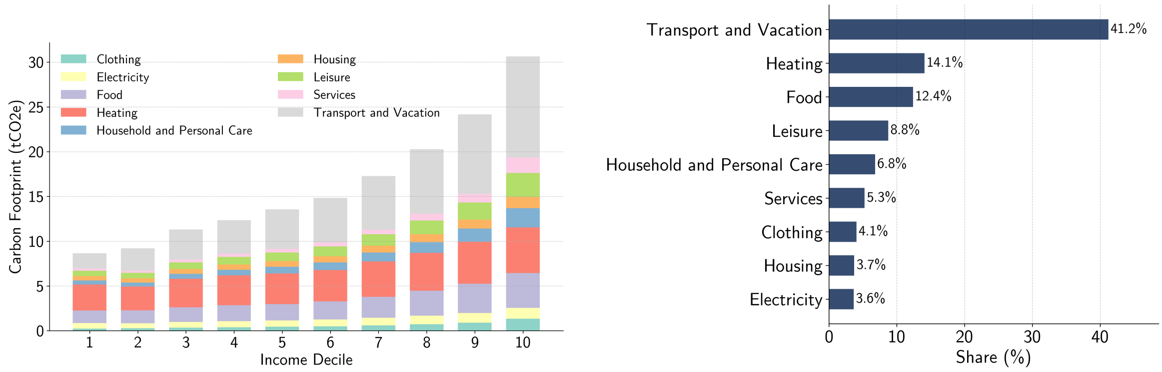

(a) Carbon Footprint by Product Category (b) Product Contribution to the p90–p10 Gap

Figure 12: Decomposition of Carbon Footprints

Notes: The left panel shows the average carbon footprint by income decile of the annual gross labor income distribution and broad product category. The right panel shows each product category's contribution to the gap in the carbon footprint between the 90th and 10th income percentile.

creasingly overstates the carbon footprint calculated based on the detailed COICOP products. Quantitatively, the aggregation error is small with an average across the income distribution of -0.68%. The overestimation at the top reflects within-category composition effects: high income tax units spend relatively more on goods within broad categories whose detailed COICOP emission intensities differ from the category-level average used in the aggregated classification.

E Quantitative Model

E.1 Initial Effective Carbon Price

We approximate the initial effective carbon price on emissions embodied in German household consumption in 2018 as the uniform carbon price that would generate the same aggregate carbon cost as the EU ETS component embodied in the 2018 consumption basket. Since only part of embodied emissions is covered by the EU ETS, this effective price is given by the 2018 average EU ETS allowance price multiplied by the emission-weighted ETS coverage share of embodied household emissions. Using an average EU ETS price of $\tau_{2018}^{ETS} = 15.96$ EUR/tCO$_2$[67] and an effective embodied-emissions coverage share of $\bar{\rho}^{ETS} = 24.5\%$, we obtain

$$\tau_0 = \tau_{2018}^{ETS} \bar{\rho}^{ETS} = 15.96 \times 0.245 = 3.91 \text{ EUR/tCO}_2e.$$

To compute the share of embodied emissions covered by the EU ETS, we construct an ETS-covered CO_2 satellite account in the EXIOBASE product classification and apply it to the 2018 input-output model. Section E.1.1 describes this construction.

[67] We compute the average EU ETS price in 2018 as the arithmetic mean of the daily closing prices for EU Allowances (EUAs), obtained from Trading Economics (2026).

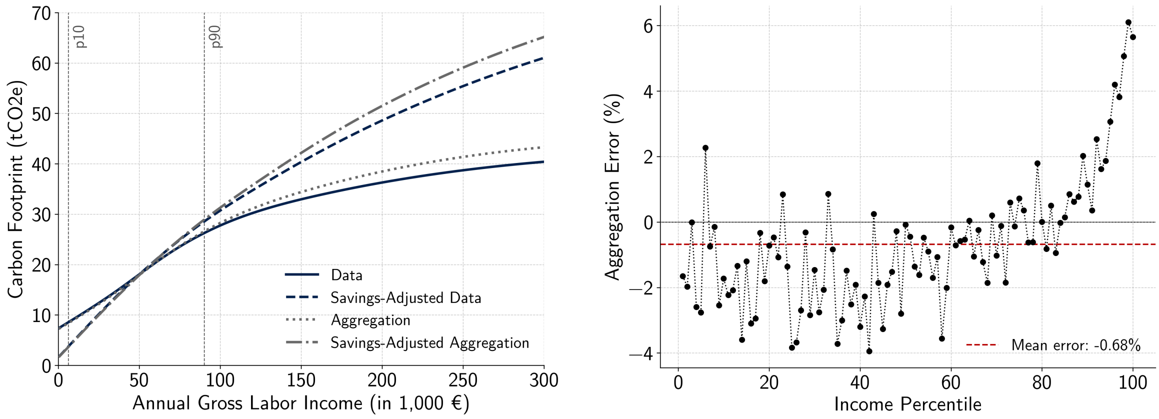

(a) Carbon Footprint (b) Aggregation Error

Figure 13: Carbon Footprint and Aggregation Error by Labor Income

Notes: The left panel shows carbon footprint estimates based on detailed COICOP categories (blue line) and on the aggregated good classification into broad products and varieties (grey line), for both the raw and savings-adjusted carbon footprints. All series show a cubic smoothing spline fit to the respective data moments. The right panel shows the aggregation error (black solid line), defined as the percentage difference between the aggregated estimate and the detailed COICOP estimate. The black dots correspond to percentile-level observations. The x-axis shows gross labor income y of a tax unit in the left panel and the gross labor income percentile in the right panel. The grey vertical dashed lines mark the 10th and 90th percentiles of the annual gross labor income distribution.

E.1.1 EU ETS Coverage

Data. We use installation-level data on verified emissions under the EU Emissions Trading System (EU ETS) compiled by Abrell (2024). The data include a match between installations and NACE Rev. 2 sectors, allowing us to aggregate verified emissions at the country and sector level.[68] Since aviation emissions are not contained in the installation-level data, we add aviation emissions from the EU ETS database (European Environment Agency 2026) and assign them to NACE Rev. 2 division 51, corresponding to air transport. The resulting dataset provides country-sector ETS emissions that can be mapped to the EXIOBASE country-product classification.

Mapping NACE sectors to EXIOBASE products. We next construct an ETS-covered CO_2 satellite account in the EXIOBASE product classification. Let M denote the number of NACE Rev. 2 sectors. We map NACE-level ETS emissions to EXIOBASE products using correspondence provided by Stadler et al. (2018). Let $\mathbf{C} \in \{0,1\}^{P \times M}$ denote the binary correspondence matrix, where $C_{pm} = 1$ if NACE sector m is mapped to EXIOBASE product p, and $C_{pm} = 0$ otherwise.

For each country k, we allocate ETS emissions from NACE sectors to the corresponding EXIOBASE products. If a NACE sector maps to multiple EXIOBASE products, we split its

[68]In 2018 around 5% of installations have no NACE code. Restricting the sample to installations with a non-missing NACE code leaves aggregate verified emissions virtually unchanged. The excluded installations account for about 1% of total verified emissions.

ETS emissions in proportion to direct CO_2 emissions in the EXIOBASE satellite account. Let $F_{km}^{ETS,NACE}$ denote ETS emissions in country k and NACE sector m, and let F_{kp}^{CO2} denote direct CO_2 emissions in EXIOBASE in country k and product p. Define the allocation weight

$$\omega_{kpm} = \frac{C_{pm} F_{kp}^{CO2}}{\sum_{q=1}^{P} C_{qm} F_{kq}^{CO2}}.$$

The ETS emissions distributed to EXIOBASE product p are then given by

$$F_{kp}^{ETS,EXIO} = \sum_{m=1}^{M} \omega_{kpm} F_{km}^{ETS,NACE} \qquad \forall\ k = 1, \ldots, K, \quad p = 1, \ldots, P.$$

The resulting matrix $\mathbf{F}^{ETS,EXIO} \in \mathbb{R}^{K \times P}$ is then vectorized, using the same country–product ordering as in the input-output system, to obtain the ETS-covered CO_2 satellite account $\mathbf{F}^{ETS} \in \mathbb{R}^{1 \times KP}$.

ETS coverage share. We compute ETS-covered indirect emission intensities by replacing the EXIOBASE emission satellite account $\mathbf{F}$ with the ETS-covered CO_2 satellite account $\mathbf{F}^{ETS} \in \mathbb{R}^{1 \times KP}$, and otherwise applying the same procedure as in Section B.3. This yields an ETS-covered indirect emission intensity f_n^{ETS} for each COICOP category n, measured in kg CO_2 per euro of household expenditure.

The ETS coverage share of embodied indirect CO_2 emissions in COICOP category n is then defined as

$$\rho_n^{ETS} = \frac{f_n^{ETS}}{f_n},$$

where f_n is the corresponding emission intensity computed from the standard EXIOBASE emission satellite account.[69] The resulting emission weighted effective ETS coverage share of embodied emissions is $\bar{\rho}^{ETS} = 24.5\%$ in 2018.

E.2 Tax-Transfer System

We estimate the parameters of the tax and transfer function separately using nonlinear least squares. The estimation sample is restricted to tax units with positive labor income. To ensure a good fit across the entire income distribution, we group observations into percentiles of the gross labor income distribution prior to estimation. The parameters are then estimated using these percentile-level data.

[69] For COICOP classes c04.5.5 (heat energy) and c04.5.1 (electricity), the raw mapped ETS coverage shares exceed one. This reflects imperfect reconciliation between verified EU ETS and EXIOBASE emissions at the country–product level, rather than true coverage above 100%. Since the purpose of the exercise is to approximate the EU ETS price embodied in household consumption, we impose the natural upper bound $\rho_n^{ETS} \leq 1$ and set the effective coverage share for these two classes equal to one.

Tax function. We estimate the parameters of the tax function $(\lambda_{tax}, \tau_{tax})$ using the following regression equation:

$$y_i^{atax} = \left(1 - \exp\left[\log(\lambda_{tax})(y_i/\bar{y})^{-2\tau_{tax}}\right]\right) y_i + u_i,$$

where y_i^{atax} denotes after-tax labor income, y_i is gross labor income of tax unit i, $\bar{y}$ is average gross labor income, and u_i is the error term.

Transfer function. We estimate the parameters of the transfer function (m, ζ) using the following regression equation

$$\mathcal{T}_{0_i} = \hat{m}\bar{y}\frac{2\exp\left[-\zeta\frac{y_i}{\bar{y}}\right]}{1 + \exp\left[-\zeta\frac{y_i}{\bar{y}}\right]} + u_i,$$

where $\mathcal{T}_{0_i}$ is the received transfer of tax unit i.

Estimation fit. Figure 14 illustrates the estimation fit by comparing estimated taxes and transfers with their data counterparts. Overall, the tax-transfer function provides a good fit of the observed tax-transfer schedule in Germany.

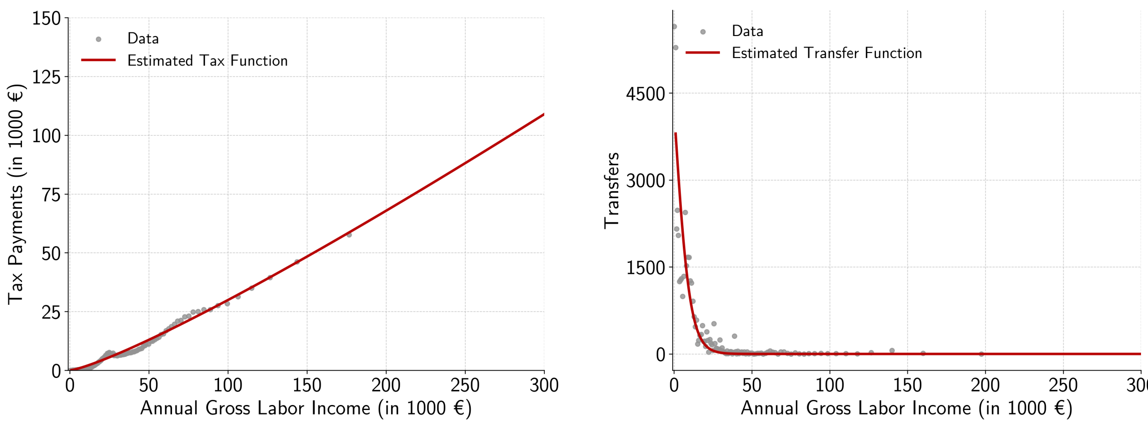

(a) Tax Function (b) Transfer Function

Figure 14: Estimated Tax-Transfer Function

Notes: The left panel shows the estimated tax function and the right panel the estimated transfer function (red solid lines). The gray dots correspond to the observed average tax payments and transfer payments, respectively. Observations are grouped into percentiles of the gross labor income distribution. The x-axis shows gross labor income y of a tax unit.

E.3 Income Distribution

Kernel density estimation. To obtain a smooth estimate of the income distribution from the LESt, we apply a standard kernel density estimation over log labor income using a Gaussian kernel and survey weights. The sample is restricted to tax units with positive labor income. We use an evenly spaced log income grid with 5000 nodes ranging from €0.8 to €10 470 000.

We select a bandwidth of approximately 0.07 on the log income scale, which implies smoothing over a local neighborhood of about 7% around each income level. Because the bandwidth is constant in log space, the corresponding smoothing window becomes wider in absolute income units as income rises.[70]

E.4 Calibration of Utility Function

E.4.1 Income Effects

Figure 15 illustrates the income effect $\eta(\theta)$ implied by the calibrated model along the gross labor income distribution. The shape and magnitude of the income effect are broadly consistent with the empirical estimates reported by Golosov, Graber, Mogstad, and Novgorodsky (2023).

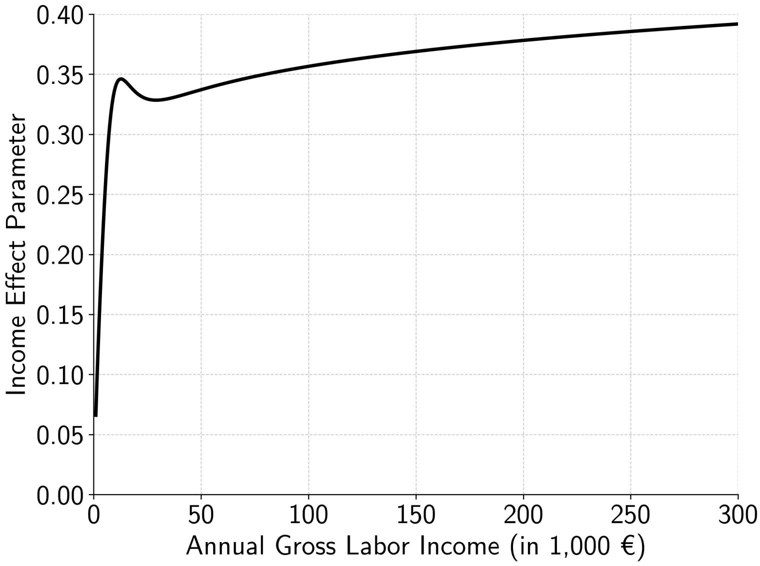

Figure 15: Income Effects

Notes: The figure shows the calibrated income effect $\eta(\theta)$. The x-axis shows gross labor income y of a tax unit.

E.4.2 Internally Calibrated Parameters

Parameters of lower-tier utility function. For each product i, we calibrate the lower-tier parameters $\boldsymbol{\theta}_i = (\epsilon_{ig}, \Omega_{ig})^\top$ by minimizing the squared distance between the model-implied and empirical within-good expenditure share on the green variety across income deciles:

$$\hat{\boldsymbol{\theta}}_i = \arg\min_{\boldsymbol{\theta}_i} \left(\mathbf{s}_{ig}(\boldsymbol{\theta}_i) - \bar{\mathbf{s}}_{ig}\right)^\top \left(\mathbf{s}_{ig}(\boldsymbol{\theta}_i) - \bar{\mathbf{s}}_{ig}\right), \tag{64}$$

[70]Using a large bandwidth is a common approach in the optimal income taxation literature, as it helps to produce smooth schedules of optimal policies (Lockwood 2020, Choné and Laroque 2010).

where $\mathbf{s}_{ig}(\boldsymbol{\theta}_i) = \big(s_{ig}(c_{i,1};\boldsymbol{\theta}_i),\ldots,s_{ig}(c_{i,D};\boldsymbol{\theta}_i)\big)^\top$ summarizes the model-implied moments for income deciles $d = 1,\ldots,D$. The model-implied within-good expenditure share on the 'green' variety in decile d is

$$s_{ig,d}(\boldsymbol{\theta}_i) = \frac{\Omega_{ig} p_{ig}^{1-\sigma_i} c_{i,d}(\boldsymbol{\theta}_i)^{\epsilon_{ig}}}{\Omega_{ib} p_{ib}^{1-\sigma_i} c_{i,d}(\boldsymbol{\theta}_i)^{\epsilon_{ib}} + \Omega_{ig} p_{ig}^{1-\sigma_i} c_{i,d}(\boldsymbol{\theta}_i)^{\epsilon_{ig}}}. \tag{65}$$

We compute the $c_{i,d}(\boldsymbol{\theta}_i)$ from the expenditure level on product i in income decile d. Let $e_{i,d} = \omega_{i,d}\bar{e}_d$ denote the model-consistent expenditure on product i in decile d, where $\bar{e}_d$ is average disposable income in decile d and $\omega_{i,d}$ is the corresponding EVS expenditure share of category i. For each candidate parameter vector $\boldsymbol{\theta}_i$, the consumption index $c_{i,d}(\boldsymbol{\theta}_i)$ is defined as the solution to

$$e_{i,d} = e_i(c_{i,d}(\boldsymbol{\theta}_i)) = \left[\sum_{j\in\{b,g\}} \Omega_{ij} p_{ij}^{1-\sigma_i} c_{i,d}(\boldsymbol{\theta}_i)^{\epsilon_{ij}}\right]^{\frac{1}{1-\sigma_i}}.$$

Substituting the resulting $c_{i,d}(\boldsymbol{\theta}_i)$ into (65) yields the model-implied expenditure share on the 'green' variety $s_{ig,d}(\boldsymbol{\theta}_i)$ used in (64). The empirical counterpart is $\bar{\mathbf{s}}_{ig} = (\bar{s}_{ig,1},\ldots,\bar{s}_{ig,D})^\top$, where $\bar{s}_{ig,d}$ is the survey-weighted average expenditure share of the brown variety in total expenditure on category i among tax units in income decile d.[71] We solve (64) by nonlinear least squares.

Parameters of upper-tier utility function. We calibrate the upper-tier parameters $\boldsymbol{\Phi} = \{\varepsilon_i, \Omega_i\}_{i=1}^{I-1}$ by minimizing the squared distance between the model-implied and empirical product expenditure shares across income deciles:

$$\hat{\boldsymbol{\Phi}} = \arg\min_{\boldsymbol{\Phi}} \ \big(\mathbf{s}(\boldsymbol{\Phi}) - \bar{\mathbf{s}}\big)^\top \big(\mathbf{s}(\boldsymbol{\Phi}) - \bar{\mathbf{s}}\big), \tag{66}$$

where $\mathbf{s}(\boldsymbol{\Phi}) = (s_{1,1}(\boldsymbol{\Phi}),\ldots,s_{I-1,1}(\boldsymbol{\Phi}),\ldots,s_{1,D}(\boldsymbol{\Phi}),\ldots,s_{I-1,D}(\boldsymbol{\Phi}))^\top$ summarizes the model-implied expenditure shares for products $i = 1,\ldots,I-1$ and income deciles $d = 1,\ldots,D$. The model implied moment for product i and income decile d is

$$s_{i,d}(\boldsymbol{\Phi}) = \frac{\Omega_i P_{i,d}^{1-\sigma} \mathbb{C}_d(\boldsymbol{\Phi})^{\boldsymbol{\epsilon_i}}}{\sum_{k=1}^{I} \Omega_k P_{k,d}^{1-\sigma} \mathbb{C}_d(\boldsymbol{\Phi})^{\epsilon_k}} \tag{67}$$

where $P_{i,d}$ denotes the price index of product i, implied by the calibrated lower-tier nest. In particular,

$$P_{i,d} \equiv \frac{e_{i,d}}{c_{i,d}} = \frac{\left[\sum_{j\in\{b,g\}} \Omega_{ij} c_{i,d}^{\epsilon_{ij}} p_{ij}^{1-\sigma_i}\right]^{\frac{1}{1-\sigma_i}}}{c_{i,d}}.$$

[71] We smooth the empirical expenditure share profiles across deciles with a locally weighted (LOESS) regression prior to calibration to limit the influence of sampling noise.

Product	Lower-Tier Utility		Upper-Tier Utility	
	Ω_{ig}	ϵ_{ig}	Ω_i	ϵ_i
Clothing	0.013	-0.983	1.000	1.000
Food	1.604	-1.022	6.016	0.620
Heating	9736.949	-1.830	32.290	0.113
Household and Personal Care	0.122	-0.940	1.604	0.967
Housing	48.742	-1.332	23.555	0.428
Leisure	15.161	-1.713	4.523	0.691
Services	9999.515	-1.868	29.722	0.596
Transport and Vacation	1.701	-0.925	4.321	1.225
Electricity	–	–	2.555	1.081

Table 8: Internally Calibrated Parameters

Notes: The table reports the internally calibrated parameters of the utility function. The moments targeted for the parameters of the lower-tier utility function are the within-good green expenditure shares; those targeted for the upper-tier utility function are the good-specific expenditure shares. The upper-tier CES weight and non-homotheticity parameter of clothing is normalized to one. Electricity is not divided into brown and green varieties because this category contains only a single product. Consequently, electricity has no lower-tier utility nest and the parameters Ω_{ig} and ϵ_{ig} need not be calibrated.

We compute the aggregate upper-tier consumption index $\mathbb{C}_d(\mathbf{\Phi})$ from the average overall expenditure E_d of each income decile d. For each candidate parameter vector $\mathbf{\Phi}$, the consumption index $\mathbb{C}_d(\mathbf{\Phi})$ is defined implicitly by the upper-tier expenditure function

$$E_d = E(\mathbb{C}_d(\mathbf{\Phi})) = \left[\sum_{k=1}^{I} \Omega_k P_{k,d}^{1-\sigma} \mathbb{C}_d(\mathbf{\Phi})^{\varepsilon_k}\right]^{\frac{1}{1-\sigma}} .$$

Substituting the resulting $\mathbb{C}_d(\mathbf{\Phi})$ into (67) yields the model-implied product expenditure shares $s_{i,d}(\mathbf{\Phi})$. The empirical counterpart is $\bar{\mathbf{s}} = (\bar{s}_{2,1}, \ldots, \bar{s}_{I,1}, \ldots, \bar{s}_{2,D}, \ldots, \bar{s}_{I,D})^{\top}$, where $\bar{s}_{i,d}$ is the survey-weighted expenditure share of product category i in total expenditure among tax units in income decile d. We solve (66) by nonlinear least squares.

Estimated parameters. Table 8 reports the calibrated parameters of the lower-tier and upper-tier utility.

Model fit. Figure 16(a) illustrates the expenditure shares of each of the nine products and Figure 16(b) shows the expenditure share of the brown variety within the corresponding category across deciles of the gross labor income distribution. The solid red lines represent the results from the model, while the dashed grey lines are the corresponding empirical moments from the EVS. The model captures both the overall levels and the distributional patterns of expenditures across income groups well.

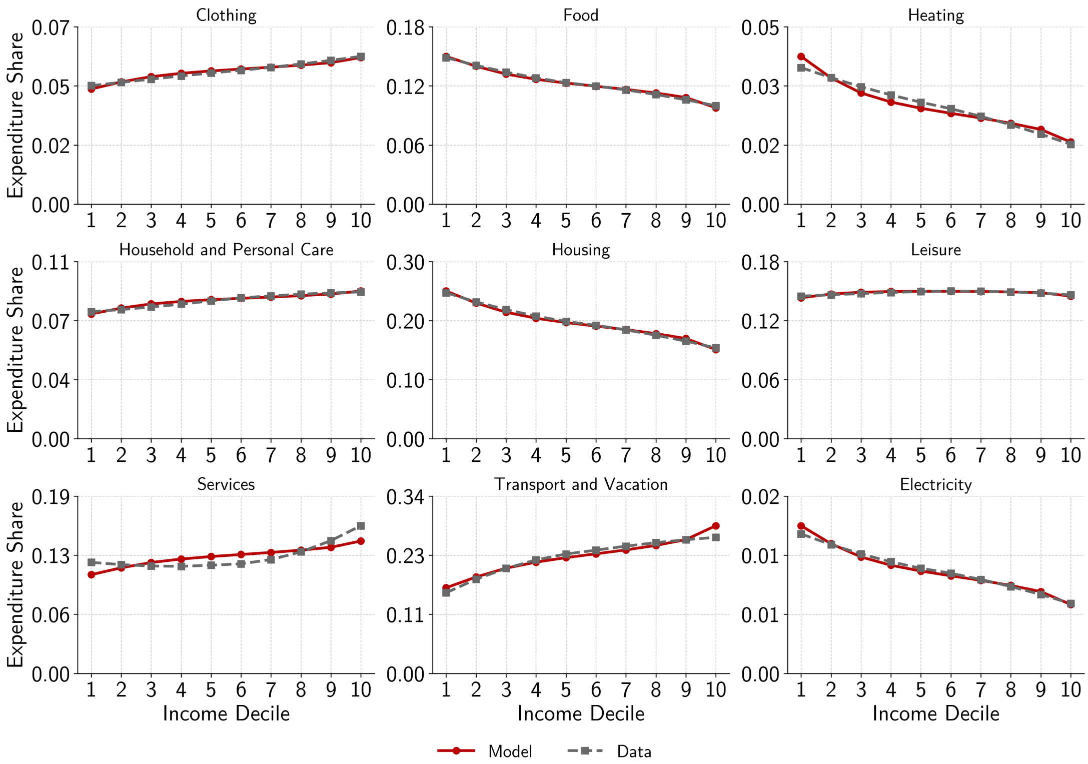

(a) Expenditure Share (Products)

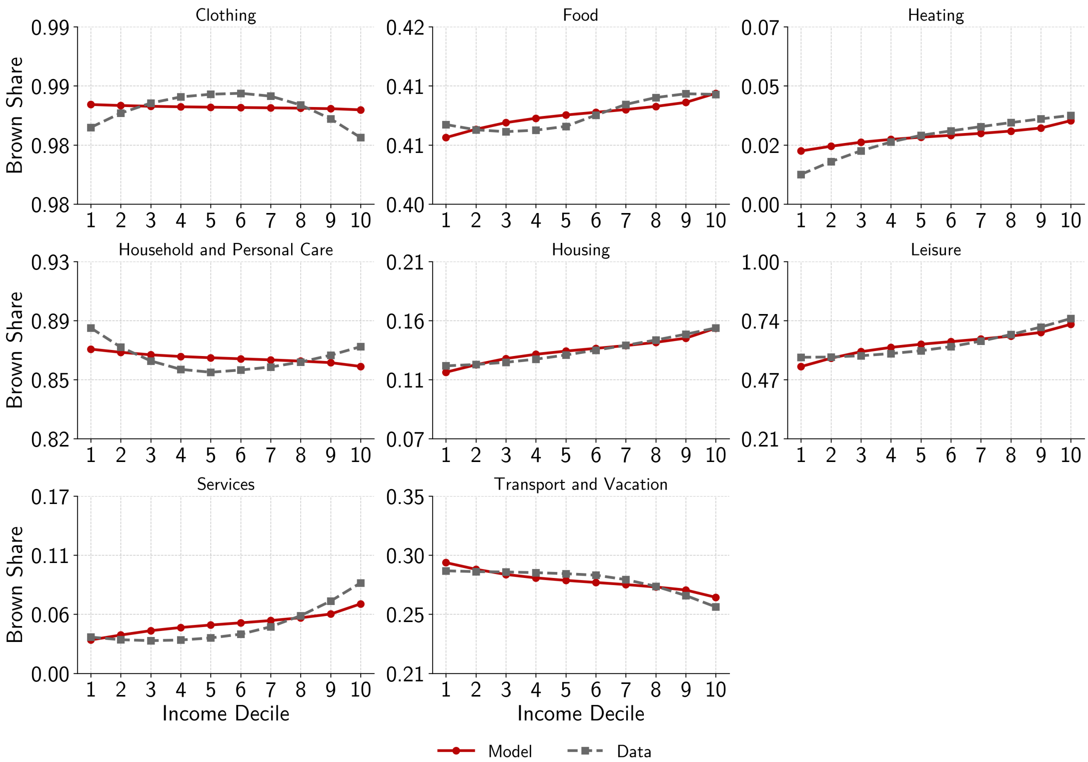

(b) Expenditure Share (Varieties)

Figure 16: Model Fit

Notes: The upper panel shows the model implied expenditure share for each of the nine products and the lower panel shows the expenditure share on the brown variety within the corresponding product category (solid line). This is compared to the dashed line which shows the expenditure shares in the EVS data (dashed line). Electricity is not divided into brown and green varieties because this category contains only a single product. Consequently, electricity has no lower-tier utility nest. The x-axis shows the decile of the annual gross labor income distribution.

E.5 Welfare Weights

Figure 17(a) shows the calibrated social marginal utility and Figure 17(b) shows the Pareto weights along the gross labor income distribution.

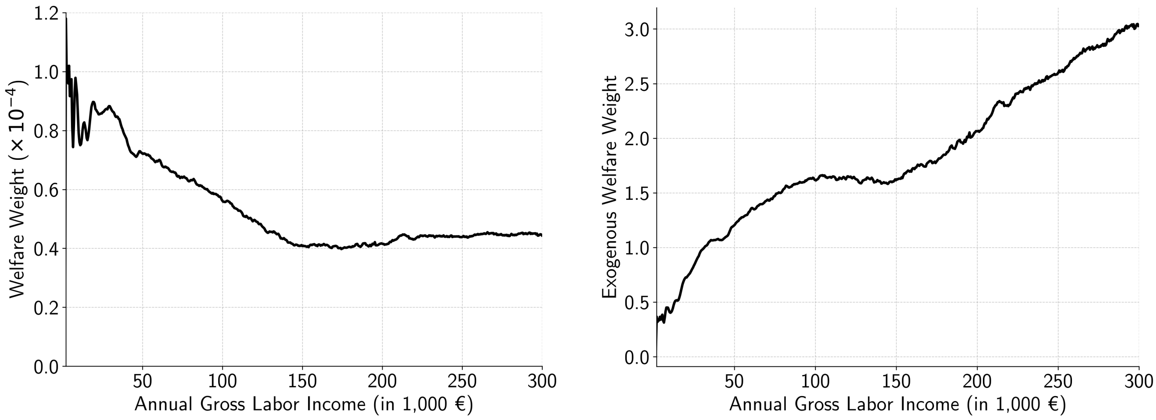

(a) Social Marginal Utility

(b) Exogenous Welfare Weights

Figure 17: Calibrated Welfare Weights

Notes: The left panel shows the social marginal utility $g(\theta, \mathcal{P}_{\text{sq}})$ and the right panel shows the Pareto weights $\omega(\theta)$ as a function of gross labor income.

E.6 Calibration of Emission Reduction Targets

This appendix describes how we map the 2030 emission targets under the German Federal Climate Change Act (*Bundes-Klimaschutzgesetz*, KSG), the EU Effort Sharing Regulation (ESR), and the EU's second emissions trading system for buildings and road transport (ETS2) into reductions relative to the model's 2018 baseline.[72]

For policy j, let q_j denote the target emissions level as a share of emissions in its statutory base year b_j, and let E_t^j denote the corresponding official emissions measure. The target emissions level relative to 2018 is then

$$\widetilde{\kappa}_j = q_j \frac{E^j_{b_j}}{E^j_{2018}},$$

implying a reduction of $1 - \widetilde{\kappa}_j$ relative to 2018. Table 9 reports the resulting mappings.

[72]Our model imposes caps on the aggregate carbon footprint generated by household consumption, whereas the statutory targets are defined in terms of territorial emissions or selected territorial source categories. The calculations should therefore be interpreted as proportional benchmarks rather than as exact equivalents to the model's emission cap.

Policy target	Baseline year	Target year	Implied 2018 emissions level	Reduction
KSG	1990	2030	$\frac{0.35 \times 1252}{852.9} \simeq 0.51$	49%
ESR	2005	2030	$\frac{0.50 \times 477.8}{434.0} \simeq 0.55$	45%
ETS2	2005	2030	$\frac{0.58 \times 344.5}{271.9} \simeq 0.73$	27%

Table 9: Mapping 2030 Emission Targets to the 2018 Baseline

Notes: KSG and ESR emissions are measured in $MtCO_2e$. Because ETS2 covers CO_2 emissions from fuel combustion in selected sectors, its mapping uses German CO_2 emissions from road transport and residential and commercial fuel combustion as an approximate measure of ETS2-covered emissions.

F Quantitative Results

	$\bar{\tau}_{\text{eff}}(y_{10})$	$\bar{\tau}_{\text{eff}}(y_{90})$	$\Pi_{90/10}$
Status quo	-9.79	30.18	39.97
Optimal carbon policy	-11.44	30.53	41.97
Carbon dividend	-16.52	30.85	47.37
Optimal carbon policy $-$ status quo	-1.64	0.36	2.00
Carbon dividend $-$ status quo	-6.73	0.67	7.40

Table 10: Change in Redistribution

Notes: The table reports the effective average tax rate (EATR) at the 10th percentile, $\bar{\tau}_{\text{eff}}(y_{10})$, and the 90th percentile, $\bar{\tau}_{\text{eff}}(y_{90})$, of the gross labor income distribution, defined as $\bar{\tau}_{\text{eff}}(\theta) = (\mathcal{T}(y(\theta)) + \tau_{co_2} f(\theta))/y(\theta)$. The EATR gap is defined as $\Pi_{90/10} = \bar{\tau}_{\text{eff}}(y_{90}) - \bar{\tau}_{\text{eff}}(y_{10})$. The final two rows report the changes relative to the status quo. EATRs are reported in percent; differences are reported in percentage points.

VAT-equivalent additional carbon-tax revenue. To benchmark the fiscal magnitude of carbon pricing, we express the increase in carbon tax revenue relative to 2018 in terms of the increase in the statutory VAT rate that would generate the same additional revenue. This comparison is purely a reporting device: VAT is not modeled explicitly in the quantitative analysis and therefore does not affect household behavior or the optimal policy.

Let τ_0 and F_{2018} denote the effective carbon price and aggregate emissions in the calibrated 2018 baseline, and let F_{cap} denote emissions under a given carbon cap. The total effective carbon price under the cap is $\tau_{\text{cap}} = \tau_0 + \Delta\tau_{co2}$, where $\Delta\tau_{co2}$ is the carbon price increase.

Let E_{cap} denote aggregate household consumption expenditure under the carbon policy, measured at purchaser prices. We construct a hypothetical net-of-VAT consumption base by applying the German standard VAT rate $t_{\text{VAT},0} = 0.19$ uniformly:

$$B_{\text{VAT}} = \frac{E_{\text{cap}}}{1 + t_{\text{VAT},0}}.$$

This calculation assumes that VAT is levied on the carbon-tax inclusive price.

Let $\Delta\mathcal{R}_{co2}$ denote the change in carbon-tax revenue relative to the calibrated 2018 baseline. We define Δt_{VAT} as the revenue-equivalent change in the statutory VAT rate. Applied to the fixed hypothetical net-of-VAT base B_{VAT}, it generates revenue equal to $\Delta\mathcal{R}_{co2}$:

$$\Delta t_{\text{VAT}} B_{\text{VAT}} = \Delta\mathcal{R}_{co2}.$$

Combining the preceding expressions gives

$$\Delta t_{\text{VAT}} = (1 + t_{\text{VAT},0})\frac{\Delta\mathcal{R}_{co2}}{E_{\text{cap}}} = (1 + t_{\text{VAT},0})\frac{\tau_{\text{cap}} F_{\text{cap}} - \tau_0 F_{2018}}{E_{\text{cap}}}.$$

The calculation abstracts from reduced VAT rates, VAT exemptions, and behavioral responses to a hypothetical VAT increase. If carbon-tax revenue under a very tight emissions cap falls below its 2018 level, the VAT-equivalent measure becomes negative, corresponding to a revenue-equivalent reduction in the VAT rate.

F.1 Along the Abatement Path

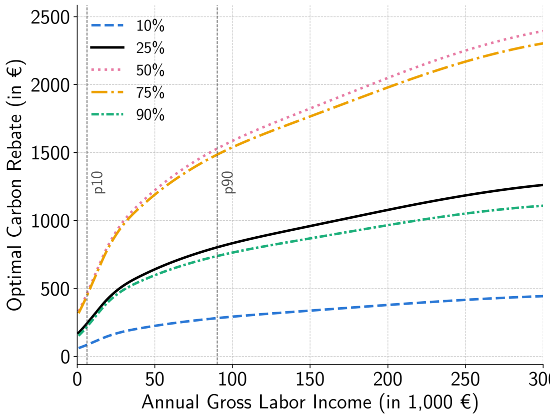

(a) Optimal Carbon Rebate

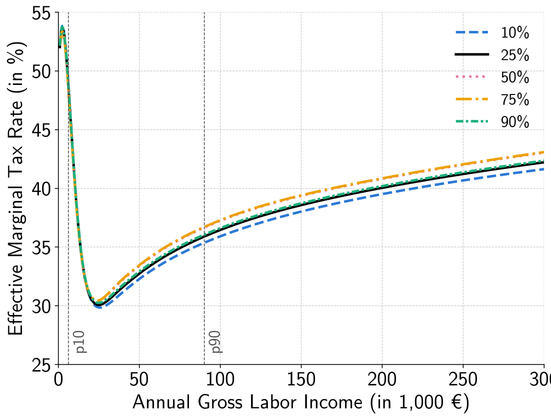

(b) Effective Marginal Tax Rates

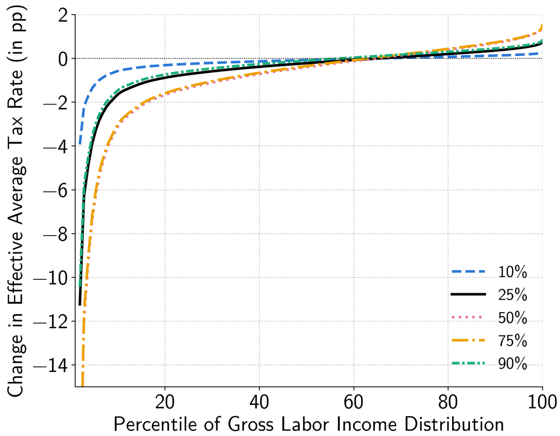

(c) Net Burden of Carbon Policy

Figure 18: Different Emission Reduction Targets

Notes: The upper left panel plots the optimal carbon rebate and the upper right panel shows the EMTR τ_{eff}, both as functions of gross labor income. The lower panel shows the change in the effective average tax rate $\bar{\tau}_{eff}$ across percentiles of the gross labor income distribution, defined as $\bar{\tau}_{\text{eff}}(\theta) = \left(\mathcal{T}(y(\theta)) + \tau_{co_2}\, f(\theta)\right)/y(\theta)$. Each curve corresponds to an aggregate emissions-reduction target of 10%, 25%, 50%, 75%, or 95% relative to the model's 2018 baseline emissions. The grey vertical dashed lines mark the 10th and 90th percentiles of the annual gross labor income distribution.

F.2 Robustness

We consider ten alternative parameter specifications, with the results summarized in Table 4. Figure 19 reports the optimal carbon rebate, EMTR, and net burden of the carbon policy along the gross labor income distribution under the alternative Frisch-elasticity specifications. Figure 20 presents the corresponding results for all remaining parameter variations. We discuss each specification in turn below.

F.2.1 Frisch Elasticity

Constant elasticity. We first vary the Frisch elasticity from its baseline value of $\varphi = 0.5$ to $\varphi = 0.25$ and $\varphi = 0.75$. Figure 19 compares the optimal carbon rebate schedule, the EMTR, and the change in the effective average tax rate along the gross labor income distribution for the baseline and the three alternative Frisch-elasticity specifications.

A lower elasticity makes the optimal rebate schedule steeper: the regressivity measure rises from 3.31 to 4.45, while the progressivity ratio of the carbon tax is almost unchanged. Consequently, the increase in redistribution of the carbon policy is smaller: the EATR gap rises by only 1.22 percentage points. A higher elasticity produces the opposite pattern: the rebate ratio falls to 2.80 and the EATR gap increases by 2.51 percentage points.

Income-varying elasticity. Finally, we allow the Frisch elasticity to increase with income from $\varphi = 0.2$ to $\varphi = 0.6$, motivated by Gruber and Saez (2002). In particular, we set $\varphi(y) = 0.2$ at and below the 10th percentile and $\varphi(y) = 0.6$ at and above the 90th percentile, with a logistic interpolation between these cutoffs. Figure 19(d) illustrates the Frisch elasticity along the gross labor income distribution.

The optimal rebate schedule becomes steeper than in the baseline, with the regressivity ratio of the carbon rebate increasing from 3.31 to 4.43. The progressivity ratio of carbon tax payments remains virtually unchanged at 7.93. The increase in the EATR gap declines to 0.90 percentage points. The reason is that the rebate at the 10th percentile falls proportionally more than the rebate at the 90th percentile, which raises the regressivity ratio and reduces the increase in the EATR gap.

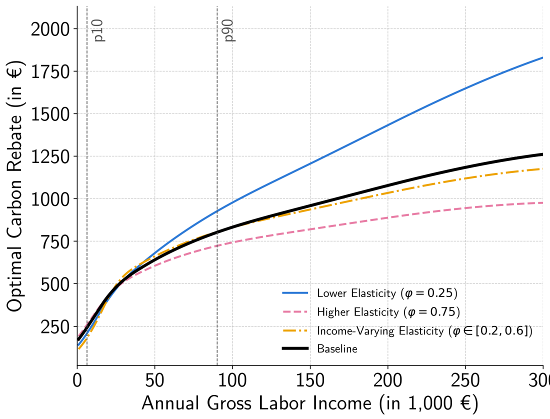

(a) Optimal Carbon Rebate

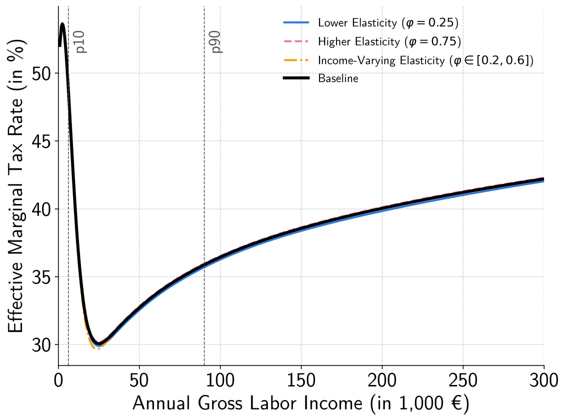

(b) Effective Marginal Tax Rates

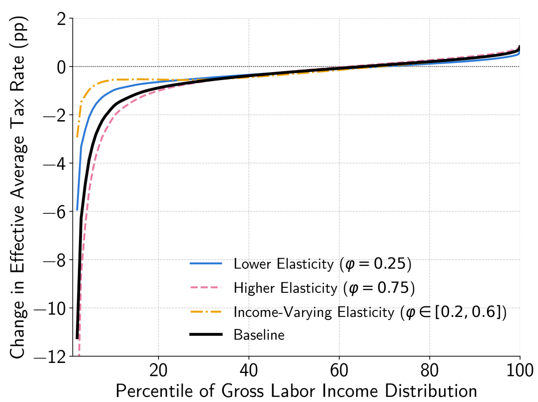

(c) Net Burden of Carbon Policy

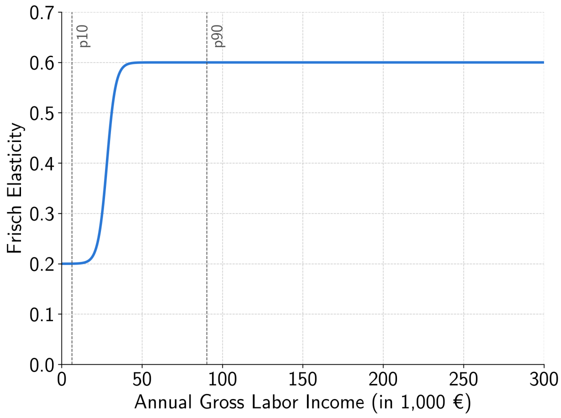

(d) Income-Varying Frisch Elasticity

Figure 19: Robustness Frisch Elasticity

Notes: The upper left panel plots the optimal carbon rebate and the upper right panel shows the EMTR τ_{eff}, both as functions of gross labor income. The lower left panel shows the change in the effective average tax rate $\bar{\tau}_{eff}$ across percentiles of the gross labor income distribution, defined as $\bar{\tau}_{\text{eff}}(\theta) = \left(\mathcal{T}(y(\theta)) + \tau_{co_2} f(\theta)\right)/y(\theta)$. Each curve corresponds to a different robustness check. The lower-right panel shows the Frisch elasticity profile used in the income-varying specification. The elasticity equals 0.2 up to the 10th percentile of the gross labor income distribution, increases smoothly according to a logistic function between the 10th and 90th percentiles, and equals 0.6 from the 90th percentile onward. The grey vertical dashed lines mark the 10th and 90th percentile of the annual gross labor income distribution.

F.2.2 Higher Curvature

In this scenario, we increase the curvature of consumption utility to $\gamma = 1.5$. The optimal rebate again becomes more regressive. As the government expands redistribution, the gains from redistribution (as captured by the distributional gains term) diminish more quickly than in the benchmark case with a lower curvature. Consequently, it is optimal to increase redistribution by less, and the EATR gap amounts to only 1.5 percentage points. The change in the curvature also affects income effects. Below, we consider an alternative scenario in which these income effects are held constant.

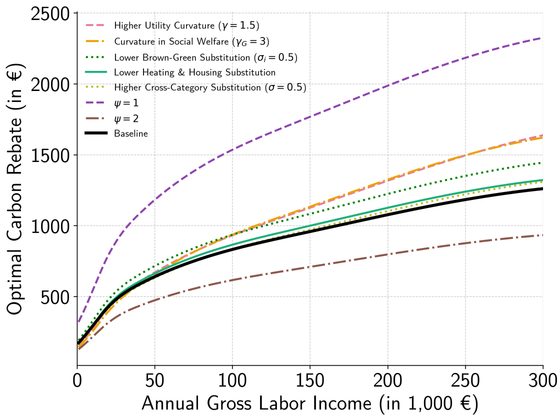

(a) Optimal Carbon Rebate

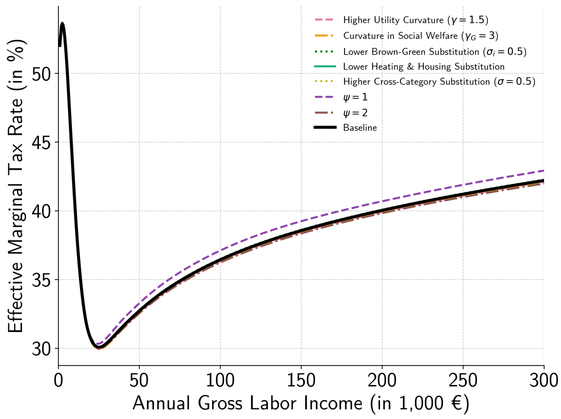

(b) Effective Marginal Tax Rates

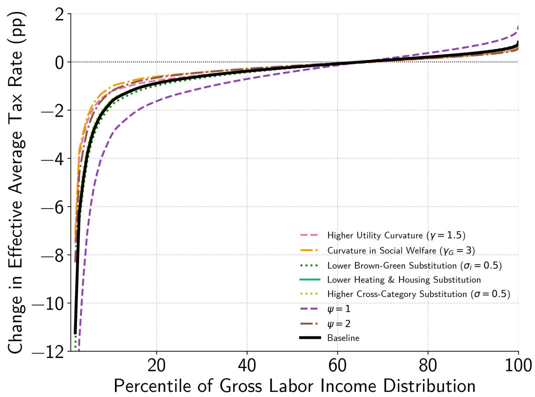

(c) Net Burden of Carbon Policy

Figure 20: Robustness Preference Parameters

Notes: The upper left panel plots the optimal carbon rebate and the upper right panel shows the EMTR τ_{eff}, both as functions of gross labor income. The lower panel shows the change in the effective average tax rate $\bar{\tau}_{eff}$ across percentiles of the gross labor income distribution, defined as $\bar{\tau}_{\text{eff}}(\theta) = \left(\mathcal{T}(y(\theta)) + \tau_{co_2} f(\theta)\right)/y(\theta)$. Each curve corresponds to a different robustness check. The grey vertical dashed lines mark the 10th and 90th percentile of the annual gross labor income distribution.

F.2.3 Elasticity of Substitution

Higher elasticity of substitution between products. We increase the elasticity of substitution across products to $\sigma = 0.5$. Greater substitutability allows households to adjust their expenditure more strongly in response to the carbon tax. This reduces the carbon price slightly. The relative incidence of the policy changes only modestly. The rebate ratio and progressivity ratio slightly increase and the EATR gap declines from 2.0 to 1.9 percentage points.

Lower elasticity of substitution between varieties. In this scenario, a higher carbon tax is required to achieve the same reduction in the aggregate carbon footprint. The regressivity of the optimal rebate and the progressivity of the carbon tax are similar to the benchmark. The increase in redistribution, $\Delta\Pi_{90/10}$, is slightly larger. However, this is primarily driven by the fact that both the level of the rebate and the carbon tax rate are higher in this scenario.

Lower substitution for housing and heating. We lower the elasticity of substitution between brown and green varieties to $\sigma_i = 0.5$ for heating and housing, while retaining $\sigma_i = 1.5$ for all other products. In these categories, substitution toward the green variety typically requires replacing the heating system, retrofitting the dwelling, or moving, rather than marginal reallocations of expenditure. Descriptive evidence supports limited substitution in the short to medium run. Roughly half of German households rent, which often limits their control over the installed heating system. Moreover, evidence from the RWI German Residential Energy Consumption Survey (GRECS) indicates that only 1.71% of non-moving households switch from a brown heating system to district heating or a heat pump over approximately five years.[73]

Because the expenditure shares on heating and housing decline with income, reducing σ_i in these categories constrains adjustment more strongly for low-income households, giving high-income households relatively greater scope to substitute toward green varieties. Nevertheless, the rebate and carbon tax incidence ratios and the implied increase in redistribution remain virtually unchanged.

F.2.4 Curvature in Welfare Function

We assume $\mathcal{G}(x) = \frac{x^{1-\gamma_{\mathcal{G}}}}{1-\gamma_{\mathcal{G}}}$ with $\gamma_{\mathcal{G}} = 3$ for the calibration of the inverse-optimum weights. This substantial increase in curvature implies a markedly smaller increase in redistribution and a rebate schedule that rises more steeply with income. The effect is similar to the second robustness check with $\gamma = 1.5$, but here the curvature effect on the social marginal utility is isolated because income effects are held constant.

F.2.5 Curvature of Abatement Costs

We vary the curvature parameter of abatement costs by setting $\psi = 1$ and $\psi = 2$. Holding the backstop price fixed, a lower (higher) value of ψ implies that firms initially abate less (more) and hence a larger carbon tax is needed to achieve the 25% reduction: it amounts to 109.40 €/tCO_2e under $\psi = 1$ and 44.69 €/tCO_2e under $\psi = 2$. Consequently, the EATR gap increases to 3.67 percentage points under $\psi = 1$, compared with 1.48 percentage points under $\psi = 2$. The relative incidence of the policy, however, is virtually unchanged: the 90–10 ratios of optimal rebates and carbon tax payments remain close to 3.3 and 7.9, respectively.

[73]We use the panel component of GRECS and link households across the 2006–08, 2008–11, and 2011–13 survey waves; see RWI and forsa (2015). To separate changes in heating systems from changes associated with relocation, we restrict the sample to households that did not move between observations. We define a switch as a transition from oil, gas, coal, wood, or stove heating in the earlier wave to district heating or a heat pump in the later wave. The corresponding transition rates are 0.84% between 2006–08 and 2008–11, 0.95% between 2008–11 and 2011–13, and 1.71% over the full interval from 2006–08 to 2011–13. Because the 2011–13 wave does not allow us to identify moves in 2009–10, the five-year transition rate is, if anything, an upper bound for households that did not move. These transition rates motivate, but do not identify, the lower substitution elasticity.