The London School of Economics and Political Science

Essays on the Macroeconomics of Climate Change and Structural Transformation

Derek Pillay

A thesis submitted to the Department of Economics
of the London School of Economics and Political Science
for the degree of Doctor of Philosophy.

London, October 2022

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 34,600 words.

Statement of conjoint work

I confirm that I am the sole author of all chapters of this thesis.

Statement of inclusion of previous work

I confirm that Chapter 3 of this thesis is a revised version of the paper I submitted at the end of my MRes in 2017.

Statement of use of third party for editorial help

I confirm that I did not receive any third party copy editing assistance with this thesis.

Abstract

The three chapters of this thesis examine climate change policy from a macroeconomic perspective, and how cross-country patterns of structural transformation are affected by international trade.

Chapter 1 studies how demographic change affects optimal carbon taxation. The model combines the standard climate-economy feedbacks from the climate economics literature with an overlapping generations structure that admits changes in fertility rates and life expectancy. I start by deriving the optimal tax-to-GDP ratio analytically, clearly illustrating the channels through which demographic change affects the optimal carbon tax. Quantitatively, I find that accounting for demographic change has a significant impact on the optimal carbon tax and leads to large welfare gains.

Chapter 1 assumes a globally harmonised climate policy that is unconstrained by the goals of the Paris Agreement. In contrast, Chapter 2 studies how carbon taxes should be set at the country-level while remaining consistent with the Paris Agreement. To ensure this consistency, each country in the model must comply with an exogenously imposed carbon budget that the policymaker is able to allocate intertemporally by setting the path of carbon taxes. I also assume that countries in my model are small and open, taking the path of climate change and interest rates as given. Using this model, I quantify the optimal path of carbon taxes and the associated output losses for a sample of 35 countries. In addition to examining the optimal policy, I assess the consequences of suboptimal carbon tax policies that provide sectoral carbon tax exemptions.

Chapter 3 examines why developing countries may be experiencing premature deindustrialisation (that is, deindustrialisation at lower levels of per capita income and with lower manufacturing shares than has historically been the norm). I start by presenting the key empirical facts, which suggest an important role for international trade. I then use a quantitative trade model to assess how the cross-country distribution of manufacturing shares is affected by various dimensions of global integration (in particular, trade in agriculture, manufacturing and services, and trade imbalances).

Dedication

To Amma and my parents, without whom this would not have been possible.

Acknowledgements

First and foremost, I would like to thank my supervisors, Professor Ricardo Reis and Professor Thomas Sampson. I am deeply indebted to them for their guidance and support throughout this process. They have greatly aided the formation of this thesis and my development as an economist.

I am also extremely grateful to all of the staff and students at the LSE who have provided me with support, encouragement and company. Special thanks go to the members of the LSE Centre for Macroeconomics for the generosity of their time and the invaluable feedback they provided, and to all the members of my PhD cohort who have shared this journey with me.

I would also like to thank my teachers, supervisors and colleagues in the UK, South Africa, Rwanda and Zanzibar who laid the foundations of my intellectual and professional development prior to my arrival at the LSE. I am especially grateful to the staff at the University of KwaZulu-Natal who first introduced me to the world of economics and who showed the utmost faith in my ability to succeed.

Many thanks go to my friends and family for their constant support and encouragement, especially my sister Clio and my uncles and aunts from Mobeni Heights.

Finally, to my wife Jess, thank you for always being patiently by my side.

Contents

List of Tables

List of Figures

Chapter 1

Optimal Carbon Taxation and Demographic Change

1.1 Introduction

The world is in the midst of a major demographic transition. For a number of decades, world life expectancy has increased while the fertility rate has declined. As a consequence, since 1970, the global population growth rate has steadily fallen while the median age has risen (see Figure 1.1). These trends are forecasted to persist at least for the duration of the twenty first century (United Nations, 2019). It has been widely recognised that this demographic shift has significant economic implications[1]. Given the intergenerational nature of climate change, it is natural to assume that accounting for demographics also has important implications for climate policy. The aim of this chapter is to investigate how long run demographic change affects one such policy: the optimal carbon tax.

Although demographic change may have implications for a wide range of policy instruments, there are a number of reasons why it is natural to start by assessing the implications for the optimal carbon tax. From a normative perspective, economists have favoured using some combination of carbon pricing and/or targeted R&D incentives as the main instruments for limiting carbon emissions[2] (Acemoglu et al., 2016; Covert et al., 2016; Nordhaus, 2007; Wall Street Journal, 2019). Moreover, carbon pricing measures have been widely implemen-

[1]Among other impacts, demographic change has been identified as an important determinant of capital flows (Krueger and Ludwig, 2007), asset prices (Carvalho et al., 2016; Lisack et al., 2021), innovation and growth (Acemoglu and Restrepo, 2022; Aksoy et al., 2019; Liang et al., 2018), and the sustainability of public finances (Attanasio et al., 2007; Poterba, 2014). For an overview of this literature, see the Handbook of the Economics of Population Aging (Piggott and Woodland, 2016).

[2]Although climate change is driven by a variety of different types of greenhouse gas (GHG) emissions, CO_2 emissions from fossil fuels are the major source of such emissions. Consequently, this chapter focuses on the regulation of carbon emissions from fossil fuels.

Figure 1.1: Global Demographic Change (1950-2100)

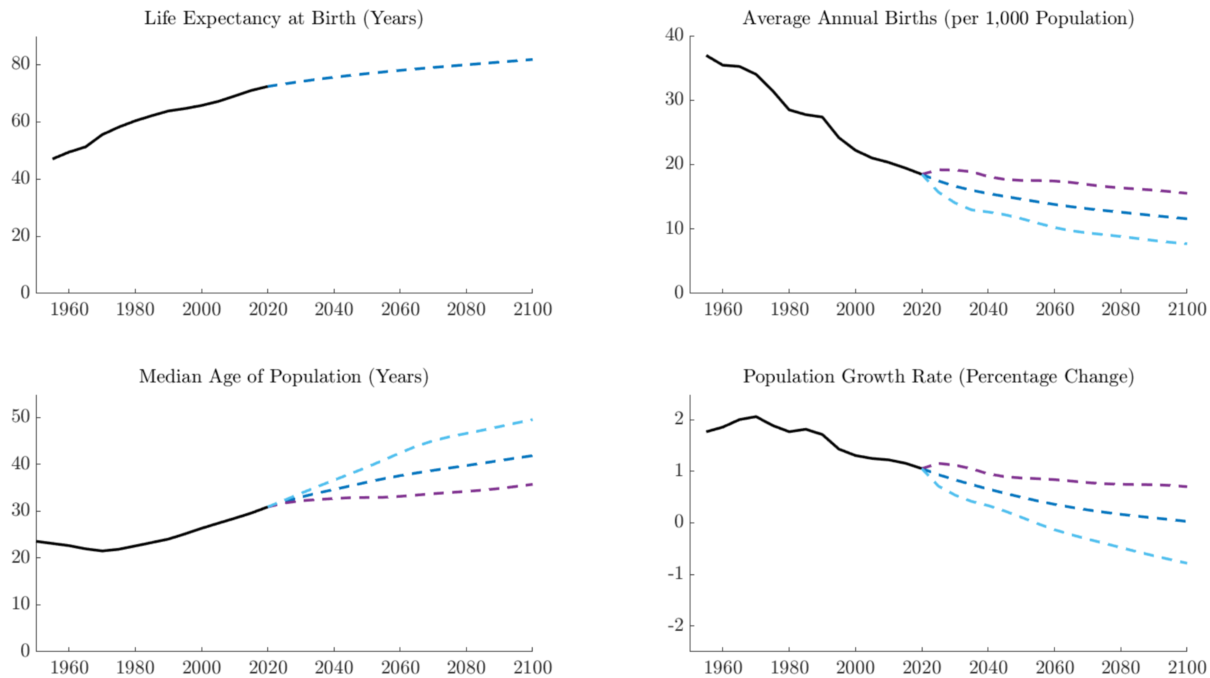

Notes: Solid black lines (1950-2020) represent historical estimates; dashed lines (2020-2100) are forecasts. The light blue, dark blue, and purple dashed lines reflect the low, medium and high projection variants, respectively.
Source: United Nations (2019).

ted in practice: as of March 2022, some form of carbon pricing was applied in 45 national jurisdictions with more to follow[3] (World Bank, 2022). In addition to these normative and positive motivations, there are strong reasons to believe that accounting for demographic change will have important implications for optimal carbon taxation. Firstly, the global population is an important determinant of the causes and effects of climate change, since humans drive climate change and suffer its consequences. It is, therefore, reasonable to assume that changes in the global population size or structure will have a direct bearing on how we should formulate carbon tax policies that make welfare tradeoffs across generations. Secondly, the general equilibrium effects of demographic change are also likely to indirectly affect the optimal carbon tax. For example, a number of authors have suggested that demographic change is responsible for the decline in real interest rates over recent decades (Carvalho et al., 2016; Rachel and Summers, 2019). It has further been argued that this fall in the real interest rate should lower the discounting of the future cost of climate change, thus implying that more aggressive climate policies in the form of higher carbon taxes are

[3]Carbon pricing typically takes two main forms - either price or quantity restrictions - which are closely related and are exactly equivalent in an environment without uncertainty or policy adjustment frictions.

optimal (Bauer and Rudebusch, 2021; Carleton and Greenstone, 2021; Gerlagh et al., 2022).

To examine these issues more carefully, I develop an integrated assessment model (IAM) with overlapping generations that captures the salient features of both climate and demographic change[4]. The model combines the standard climate-economy feedbacks from the IAM literature with an overlapping generations structure in which agents live for up to three periods: childhood, middle age and retirement. Since only a fraction of each generation transitions to the next stage of life, the model is able to accommodate the changes in fertility rates and life expectancy observed in the data and highlights the implications for optimal carbon taxation in a transparent way.

The optimal carbon tax is set to internalise the present discounted value of current and future climate-related damages that are caused by an additional unit of emissions today, accounting for any interacting market failures. In my model, demographic change influences the optimal carbon tax through four channels. First, the cost of climate change is increasing in aggregate output. Aggregate output in turn depends on the size of the labour force and the capital stock, both of which are functions of demographic parameters. Second, changes in the population size and structure over time influence the discount rate used by the policymaker[5]. Loosely speaking, more weight is placed on periods with larger populations. Due to differences in the intergenerational and intragenerational pure rates of time preference[6], however, the relevant measure of the aggregate population size is one in which each generation is weighed differently. Third, the policymaker's discount rate is also affected by changes in agents' marginal utility of consumption. The marginal utility of consumption fluctuates over time in part due to fluctuations in savings rates and aggregate output induced by demographic shifts. Fourth, the policymaker has preferences over intergenerational equity that are not internalised by the market. When the policymaker lacks the additional instruments needed to achieve the optimal distribution of consumption across agents, the optimal carbon

[4]Integrated assessment models (IAMs) are models that integrate the economy and climate in a unified framework. There are a wide variety of such models, but those most commonly employed by macroeconomists combine standard microfounded macroeconomic models with (a) a link between economic activity and the flow of carbon emissions into the atmosphere, (b) a representation (often referred to as the carbon cycle) of how the flow of emissions accumulates into a stock of atmospheric carbon concentration and (c) a representation (called a damage function) of how this stock affects economic activity and/or welfare directly. For more details, see Nordhaus (2008) and Hassler et al. (2016).

[5]To avoid ambiguity, when I refer simply to "discounting" or the "discount rate", I mean the rate used to compare the lifetime utility value of units of consumption available at different points in time. In contrast, when I refer to the "pure rate of time preference" (or "pure discount factor"), I mean the rate (or factor) used to compare the lifetime utility value of instantaneous utility enjoyed at different periods in time.

[6]The pure rate of intragenerational time preference refers to the pure rate of time preference each individual has over instantaneous utility enjoyed over the different periods of their life. The pure rate of intergenerational time preference refers to the policymaker's pure rate of time preference over the welfare of individuals born in different periods.

tax will be second best. Relative welfare across generations will be affected by demographic changes via agents' savings rates and factor returns, which interact with the policymaker's distributional motivations in setting the carbon tax.

The key finding from the baseline calibration of the model is that accounting for demographic change raises the current level of the optimal carbon tax but reduces the rate at which this tax should grow over time. In particular, the optimal carbon tax in the 2010-2040 period of the model is \$66 per ton of CO_2 (or \$243 per ton of carbon) when accounting for demographic change compared to an optimal carbon tax of \$61 per ton of CO_2 (or \$222 per ton of carbon) in a similarly calibrated representative agent setting, where both values are express in 2015 US Dollar terms[7]. That is, accounting for baseline forecasts of demographic change increases the current estimate of the optimal tax by 9.5%. The gap between the optimal tax rate with and without demographic change narrows to 1.3% in the 2040-2070 period under the baseline calibration and then from 2070 onwards, the optimal tax is essentially equivalent in both versions of the model (see Figure 1.4). These results are driven primarily by the effect of the population growth rate on the discount rate used by the planner. This result is intuitive: population growth implies that, all else equal, more weight should be placed on the future than the present. This effect partially offsets the discounting of the future due to pure rates of time preference and consumption growth present in representative agent IAMs. Since most benefits of carbon taxation accrue in the future while the costs are borne in the present, discounting the future less heavily implies that a higher carbon tax should be set at any given point in time. As the demographic shift slows, however, and the population size and structure stabilises, the optimal tax converges to the one implied by a similarly calibrated representative agent model. This convergence means that the optimal carbon tax with overlapping generations must grow less quickly than under the equivalent representative agent framework. The fact that the discrepancies between my OLG model and an equivalent representative agent framework are driven by the effect of population growth on the policymaker's discount rate is robust to the specification and calibration of the model conditional on the functional forms chosen for the utility and damage functions. Because marginal damages from emissions are proportional to aggregate output while agents' marginal utility is inversely proportional to aggregate output, discounted future damages are not a function of future aggregate output. This means that demographic changes cannot affect the optimal carbon tax via aggregate output changes. In addition, although demographic fluctuations do impact the savings rate/marginal propensity to consume and thus future marginal utility, for any reasonable parameter values, these fluctuations are small due to

[7]Unless stated otherwise, it should be assumed that all US Dollar values reported in this chapter are in 2015 US Dollar terms.

the non-linear nature of the function that maps demographic parameters into the marginal propensity to consume. The second best adjustments to the optimal carbon tax also tend to be small and diminish over time. Under the baseline calibration, the intergenerational and intragenerational pure rates of time preference are equalised. Suboptimal capital accumulation (through the eyes of the policymaker) is, therefore, driven purely by the OLG structure of the model. What matters for the optimal carbon tax is relative changes in this suboptimality over time, since the carbon tax can partially offset this suboptimality through intertemporal redistributions of output. As life expectancy and the population size converge to their constant long-run levels, the degree of suboptimal capital accumulation also becomes constant and no further second-best adjustments need to be made. In sum, this means that the channel that operates directly via the adjusted population growth rate is the dominant one, as the discounted damages in a given period are proportional to this growth rate.

The welfare effects of these carbon tax adjustments are significant. The increase in social welfare from using the optimal carbon tax formula that accounts for demographic change instead of the suboptimal representative agent formula is \$148 Billion (\$2015 lump-sum aggregate consumption equivalent). These findings highlight the importance of properly accounting for demographic change in IAMs that produce quantitative assessments of optimal climate policy. Related to this point, the findings also illustrate that a key source of uncertainty for policymakers relates to future demographic trends and in particular, assumptions about the future path of the aggregate population size. In contrast to the above results from the baseline calibration of the model, which uses the UN's medium variant population projections, a model calibration using the low and high variant population projections delivers very different results. Relative to the representative agent benchmark, the low variant projection delivers an optimal carbon tax that is 9.3% lower in 2010-2040 and 16.5% lower in 2040-2070, with the optimal tax only converging to the representative agent benchmark in 2160-2190. Conversely, the model calibrated to the high variant projection delivers an optimal carbon tax that is 36% higher that the representative agent benchmark in 2010-2040, with the gap only converging to zero in 2160-2190 (see Figure 1.4). The welfare gains of using the optimal carbon tax relative to the suboptimal representative agent formula also vary substantially with the population projections: the welfare gains under the low and high variant projections increase to \$626 Billion and \$2.2 Trillion, respectively (\$2015 lump-sum aggregate consumption equivalent). Much of the existing work around optimal climate policy and uncertainty has focused on climate and economic risks such as climate tipping points or productivity shocks within representative agent models (Brock and Hansen, 2018; Cai and Lontzek, 2019; Lemoine and Traeger, 2014; Weitzman, 2009). The sensitivity of the optimal carbon tax to different demographic scenarios suggests that accounting for demographic uncertainty should also be a key concern in formulating climate policy.

Related Literature. The work on climate-economy interactions within economics was pioneered by Nordhaus, whose DICE (1993; 1994) and RICE (Nordhaus and Yang, 1996) models were part of the first generation of IAMs[8]. IAMs have subsequently been extended in two key ways that are of relevance to my chapter. First, there are a number of papers that derive closed-form expressions for the optimal carbon tax, with these optimal carbon tax formulae often depending on a small set of parameters (Bretschger and Karydas, 2018; Dietz and Venmans, 2019; Gerlagh et al., 2022; Gerlagh and Liski, 2018; Lemoine and Rudik, 2017; Rezai and van der Ploeg, 2016; Traeger, 2015; Van den Bijgaart et al., 2016). Of these papers, this chapter is most closely related to the work of Golosov et al. (2014) as the optimal tax formula I derive nests theirs as a special case. This literature has proved valuable both in making some of the key mechanisms behind optimal carbon taxation more transparent and also in aiding quantification of the optimal carbon tax by restricting the data required to calibrate these models.

Second, there is a literature that analyses the implications of moving away from representative agent IAMs to understand the interplay between demographics and climate policy. Early contributions to this literature were made by Howarth (1996; 1998; 2000) who analysed the conditions under which OLG and representative agent IAMs coincide in their policy implications. More recent examples of such work include Gerlagh and van der Zwaan (2001); Gerlagh et al. (2022); Karp and Rezai (2014); Kotlikoff et al. (2021); Quaas and Bröcker (2016); von Below et al. (2013); Williams et al. (2015).

I contribute to the intersection of these two extensions of the IAM literature by developing an OLG IAM where the optimal carbon tax to GDP ratio can be characterised in closed form using a limited set of parameters and exogenous variables. In doing so, I am able to provide a simple and transparent quantification of how two long-run demographic trends (increasing life expectancy and decreasing fertility rates) affect the optimal carbon tax. My chapter is almost identical to Gerlagh et al. (2022) in its motivation and model structure; our results, however, are very different. While both Gerlagh et al. (2022) and myself find that accounting for demographic change raises the current level of the optimal carbon tax under baseline forecasts of demographic change, they find that moving away from a representative agent framework also significantly raises the growth rate of the optimal carbon tax. This implies a large divergence over time between our estimates of the optimal carbon tax, given that my estimates of the optimal carbon tax converge with those from the representative agent setting. Our results deviate for two reasons. First, by assuming that the government

[8]See Nordhaus and Boyer (2003) and Nordhaus (2008) for detailed and updated expositions of these models with a wide range of applications.

makes the necessary lump sum transfers across generations, Gerlagh et al. (2022) derive a first-best optimal carbon tax. More importantly, we use different formulae for the optimal carbon tax. Gerlagh et al. (2022) start with a representation of the optimal carbon tax from the corresponding representative agent setting before quantifying this formula using their calibrated OLG model. Instead, I derive the optimal carbon tax formula from first principles within the context of my framework, producing an optimal carbon tax formula and quantified results that are microfounded and internally consistent.

An additional benefit of using an overlapping generations structure is that I can address important debates in the climate economics literature about the appropriate discount rate in optimal carbon tax calculations. In particular, there is a long-standing dispute, most notably between William Nordhaus and Nicholas Stern, regarding the correct pure rate of time preference that should be applied to optimal carbon tax calculations. On one hand, Nordhaus makes a positive argument for a pure rate of time preference informed by observed asset returns and consumption growth (Nordhaus, 2008). In contrast, Stern takes a normative approach in arguing that we have little ethical reason to treat generations differently, implying a much lower pure rate of time preference that is closer to zero (Stern, 2007). Given the longterm impacts of climate policy, the optimal carbon tax is very sensitive to the assumed pure rate of time preference and this debate has been one of the most quantitatively important sources of disagreement in the literature. Through the lens of my model, Nordhaus and Stern are focusing on two separate parameters. Absent bequest motives, observables can be mapped onto individuals' pure rate of time preference exclusively over their own lives. These observables are not informative, however, of how policymakers should value welfare across generations. By introducing two pure time preference parameters in the policymaker's objective function (an intragenerational pure rate of time preference and an intergenerational pure rate of time preference) my model is able to simultaneously accommodate both the Nordhaus and Stern world views. This chapter is, therefore, also related to the climate economics literature that analyses the implications of differentiating between private and social discount rates (Barrage, 2018; Belfiori, 2017, 2018; Goulder and Williams, 2012; Kaplow et al., 2010; Mier et al., 2021; von Below, 2012). In particular, I test the sensitivity of optimal carbon tax to the two pure rates of time preference in my model. As with Quaas and Bröcker (2016), I find that the optimal carbon tax is much more sensitive to assumptions on the intergenerational rate than the intragenerational rate, implying that it is difficult to derive climate policy recommendations that are robust to normative judgements.

Furthermore, the debate about discounting has recently been revisited in light of the decline in real interest rates that occurred following the Great Recession. Real interest rates on safe assets are now significantly lower than they were when the Nordhaus-Stern debate about discounting first emerged, with a number of authors arguing that optimal carbon tax estimates should be revised upwards quite substantially (Bauer and Rudebusch, 2021; Carleton and Greenstone, 2021)[9]. My model shows that while these arguments are intuitively appealing, they are potentially misleading. Since both the real interest rate and the optimal carbon tax are endogenous objects that are simultaneously determined by the primitives of the model, it is not sufficient to simply revise the discount rates used in the optimal carbon tax formula while holding everything else constant. The underlying causes of the real interest rate decline need to be modelled and the equilibrium implications for the optimal carbon tax determined in a consistent manner. The real rate does indeed decline significantly over time in my calibrated model as a result of the two demographic shifts I focus on. First, increased life expectancy raises savings rates leading to a larger future capital stock; second, lower fertility rates reduce the number of workers in the future. Both of these effects lead to an increase in the capital-to-labour ratio, which pushes down returns[10]. While it is true that this real interest rate decline occurs while the optimal carbon tax is also rising in my model, this negative correlation is driven by the fact that my optimal carbon tax is a non-stationary object, since it converges to an object that is proportional to GDP. When analysing the optimal tax-to-GDP ratio, which is a stationary object, my model shows that the demographic transition simultaneously lowers the real interest rate and the optimal tax-to-GDP ratio. This implies that a decline in real interest rates caused by the demographic transition captured in my model should, in some sense, be associated with less aggressive climate policy than would have been the case absent the real rate decline.

Roadmap. The rest of the chapter is organised as follows. Section 2 outlines the model. Section 3 characterises the optimal carbon tax. The policymaker must take the competitive equilibrium constraints as given when choosing the carbon tax. Because the policymaker has one instrument (the carbon tax) but two goals (internalising the externality and redistributing income), the policy choice is second best. Section 4 of the chapter provides the quantitative results for a calibrated version of the model. In addition to mapping out the path of the optimal carbon tax, the implications for the optimal path of emissions, warming and climate damages are examined along with the welfare gains of implementing the optimal policy. Section 5 uses the model to address the debates in the literature regarding

[9]Strictly speaking, the papers mentioned here address the Social Cost of Carbon (SCC) rather than the optimal carbon tax. The SCC is equal to the discounted value of all current and future externality damages associated with emitting a unit of carbon today. The two are closely related, however, in that the optimal carbon tax is simply the SCC when the climate externality is the only market failure.

[10]As there are no financial frictions in my model and I calibrate the model assuming all agents have perfect foresight, the return on capital is equal in equilibrium to the real interest rate on a safe asset (such as government bonds) in zero net supply.

discounting. Section 6 concludes.

1.2 Model

This section outlines the model used in this chapter. The model combines a climate-economy structure based on Golosov et al. (2014) with an overlapping generations structure inspired by Gertler (1999) and Cipriani (2014). As in much of the IAM literature, the model is global by nature and thus abstracts from potential frictions and distributional issues across regions. There are two sectors of the economy: a final goods sector, which provides output for consumption and investment, and an energy sector, which produces an intermediate good used by the final goods sector. Within the energy sector, there are dirty and green energy varieties. Consumption of dirty energy produces carbon emissions that accumulate in the atmosphere and depreciate slowly over time. As the atmospheric carbon concentration increases and the climate changes, this leads to output losses in the final goods sector in a manner analogous to a reduction in total factor productivity (TFP). Firms take the path of climate change as given when choosing their energy input, leading to a negative externality and an inefficiently high level of emissions in the laissez-faire equilibrium. The aim of the policymaker is to set a carbon tax on dirty energy use to maximise social welfare.

1.2.1 Households

Agents live for up to three periods. Children born in period t transition to middle age with an exogenous probability given by ψ_{t+1}^{t}. The middle aged in period $t+1$ then transition to retirement with an exogenous probability given by ψ_{t+2}^{t}. Letting N_t^s denote the population alive in period t who were born in period s, this implies that

$$N_{t+1}^{t} = \psi_{t+1}^{t} N_t^t \tag{1.1}$$

$$N_{t+2}^{t} = \psi_{t+2}^{t} N_{t+1}^t \tag{1.2}$$

The young population grows exogenously at a net growth rate of n_t such that

$$N_t^t = (1 + n_t)\, N_{t-1}^{t-1} \tag{1.3}$$

The total population in period t is given by $N_t \equiv N_t^t + N_t^{t-1} + N_t^{t-2}$. The expected lifetime utility of an agent born in period t is given by

$$\mathbb{E}_t \left[U_t^t\right] = \ln c_t^t + \psi_{t+1}^{t} \beta \left(\mathbb{E}_t \left[\ln c_{t+1}^t\right] + \psi_{t+2}^{t} \beta \mathbb{E}_t \left[\ln c_{t+2}^t\right]\right) \tag{1.4}$$

where c_s^t denotes the period s consumption of an agent born in period t, β is the agent's pure discount factor and $\mathbb{E}_t$ is the expectation conditional on the time t information set. It is assumed that an agent born in period t knows the transition probabilities ψ_{t+1}^t and ψ_{t+2}^t at time t. Agents are passive when young: they supply no labour and their consumption is proportional to that of their parent, which they have no control over. The middle aged inherit no assets and each supply one unit of labour inelastically. They use their labour income to finance consumption for themselves and their children. In addition, they can invest in capital that they rent to firms during retirement with any residual firm profits also paid to retirees as dividends. Once retired, agents have no labour endowment and finance their consumption from their savings leaving no bequests for subsequent generations. The period budget constraints of the individual born in period t are therefore given by:

$$c_t^t = \theta_t^y c_t^{t-1} \tag{1.5}$$

$$c_{h,t+1} + k_{t+2}^t = w_{t+1} + d_{t+1}^t \tag{1.6}$$

$$c_{t+2}^t = \frac{R_{t+2}}{\psi_{t+2}^t} k_{t+2}^t + \frac{\Pi_{t+2}}{N_{t+2}^t} + d_{t+2}^t \tag{1.7}$$

where θ_t^y is the exogenously determined factor of proportionality between the consumption of parents and their children, $c_{h,t+1}$ is the consumption of the whole middle-aged household, w_{t+1} is the wage, k_{t+2}^t is investment in the period $t+2$ capital stock, R_{t+2} is the gross return on capital net of depreciation, Π_{t+2} are total profits in the economy, and d_s^t denotes the period s lump sum transfers made to agents born in period t. The $\frac{R_{t+2}}{\psi_{t+2}^t}$ term in (1.7) reflects the fact that there are perfect annuity markets through which the assets of those who die before retirement are redistributed to all surviving retirees. Capital fully depreciates after use and thus has no resale value, and consumption choices and capital holdings cannot be negative.

1.2.2 Firms

Final goods production uses energy, capital and labour as inputs to a Cobb-Douglas production function:

$$Y_t = \left[1 - D_t\left(S_t\right)\right] A_{y,t} E_{y,t}^{\nu} K_{y,t}^{\alpha} L_{y,t}^{1-\nu-\alpha} \tag{1.8}$$

$D_t\left(S_t\right)$ is a damage function reflecting the impact of S_t (the atmospheric carbon concentration) on the economy. The damage function takes the form

$$1 - D_t\left(S_t\right) = e^{-\gamma_t\left(S_t - \bar{S}\right)} \tag{1.9}$$

where $\bar{S}$ is the pre-industrial atmospheric carbon concentration. TFP in the final goods sector ($A_{y,t}$) evolves exogenously. $K_{y,t}$ and $L_{y,t}$ refer to the capital and labour input, respectively, of the final goods sector while $E_{y,t}$ is the use of an energy composite.

The energy composite is supplied using a CES aggregator of dirty energy ($E_{d,t}$) and green energy ($E_{g,t}$)

$$E_{y,t} = \left(\omega_d E_{d,t}^{\frac{\sigma_e - 1}{\sigma_e}} + \omega_g E_{g,t}^{\frac{\sigma_e - 1}{\sigma_e}}\right)^{\frac{\sigma_e - 1}{\sigma_e}} \tag{1.10}$$

The dirty and green energy varieties are produced linearly from labour:

$$E_{k,t} = A_{k,t} L_{k,t} \quad k \in \{d, g\} \tag{1.11}$$

where $L_{k,t}$ is labour input and $A_{k,t}$ is exogenously determined productivity.

The labour, capital, final goods and energy market clearing conditions, respectively, are given by

$$L_{y,t} + L_{d,t} + L_{g,t} = N_t^{t-1} \tag{1.12}$$

$$K_{y,t} = k_t^{t-2} N_{t-1}^{t-2} \tag{1.13}$$

$$c_t^t N_t^t + c_t^{t-1} N_t^{t-1} + c_t^{t-2} N_t^{t-2} + k_{t+1}^{t-1} N_t^{t-1} = Y_t \tag{1.14}$$

$$E_{y,t} = \left[\omega_d E_{d,t}^{\frac{\sigma_e - 1}{\sigma_e}} + \omega_g E_{g,t}^{\frac{\sigma_e - 1}{\sigma_e}}\right]^{\frac{\sigma_e - 1}{\sigma_e}} \tag{1.15}$$

$$E_{k,t} = A_{k,t} L_{k,t} \quad k \in \{d, g\} \tag{1.16}$$

1.2.3 Government

The government can levy carbon taxes (τ_t) on the use of dirty energy. The revenue from this tax is rebated lump sum to middle- and old-aged households such that the government runs a balanced budget in each period:

$$\tau_t E_t = d_t^{t-1} N_t^{t-1} + d_t^{t-2} N_t^{t-2} \tag{1.17}$$

where these transfers are restricted such that $d_t^s \geq 0$.

1.2.4 Carbon Cycle

Dirty energy use produces carbon emissions that are released into the atmosphere where they depreciate slowly over time. The atmospheric carbon concentration evolves according to

$$S_t - \bar{S} = \sum_{i=0}^{t} (1 - d_i) E_{d,t-i} \tag{1.18}$$

where $1 - d_i$ represents the amount of carbon released today that would be left in the atmosphere i periods into the future, $E_{d,t}$ is carbon emissions in period t, and period 0 in the model is taken to be the first period in which there were industrial emissions[11]. The functional form for atmospheric carbon depreciation is given by

$$1 - d_i = \varphi_L + (1 - \varphi_L) \varphi_0 (1 - \varphi)^i \tag{1.19}$$

where $\varphi_L \in [0, 1]$ is the share of carbon that will stay in the atmosphere indefinitely and $1 - \varphi_0 \in [0, 1]$ is the share of transient emissions that will exit the atmosphere immediately. The remaining share of transient emissions then decay geometrically at rate φ.

1.3 Optimal Carbon Taxation

This section first outlines the competitive equilibrium conditions of the economy before solving the laissez-faire equilibrium in which carbon taxes are set to zero. Due to the carbon externality, too much dirty energy is produced and consumed in this equilibrium leading to suboptimal outcomes. A second best carbon tax formula is then derived under the assumption that the policymaker cannot use carbon tax revenues to alter distributional outcomes contemporaneously and has no other policy instruments at their disposal. This assumption is made for analytical tractability but is also a reasonable one given how many carbon tax systems are designed in practice. The policymaker is still able to use climate policy to alter distributional outcomes intertemporally, however, which means that the optimal carbon tax is not set to simply internalise the carbon externality but is also set with the distribution of intergenerational welfare in mind.

1.3.1 Competitive Equilibrium

A competitive equilibrium in this economy is formally defined as follows

Definition 1: A competitive equilibrium consists of a sequence of allocations $\left\{c_t^t, c_t^{t-1}, c_t^{t-2}, k_{t+1}^{t-1}, E_{y,t}, K_{y,t}, L_{y,t}, E_{d,t}, E_{g,t}, L_{d,t}, L_{g,t}, S_t\right\}$, prices $\{p_{e,t}, R_t, w_t, p_{d,t}, p_{g,t}\}$ and policies $\left\{\tau_t, d_t^{t-1}, d_t^{t-2}\right\}$ such that:

[11] The units of dirty energy are normalised such that one unit of dirty energy leads to one unit of emissions.

1. the allocations solve the consumers' utility maximisation problems and the firms' profit maximisation problems given prices, polices and the climate,
2. the government budget constraint is satisfied in every period,
3. atmospheric carbon concentration satisfies the carbon cycle constraint in every period, and
4. markets clear.

The profit maximising first order conditions for the final goods firm are

$$\nu \frac{Y_t}{E_{y,t}} = p_{e,t} \tag{1.20}$$

$$\alpha \frac{Y_t}{K_{y,t}} = R_t \tag{1.21}$$

$$(1 - \nu - \alpha) \frac{Y_t}{L_{y,t}} = w_t \tag{1.22}$$

where $p_{e,t}$ is the price paid for the energy composite. Similarly, the profit maximising first order conditions for the firm producing the energy composite are

$$p_{e,t} E_{y,t}^{\frac{1}{\sigma_e}} \omega_d E_{d,t}^{\frac{-1}{\sigma_e}} = p_{d,t} + \tau_t \tag{1.23}$$

$$p_{e,t} E_{y,t}^{\frac{1}{\sigma_e}} \omega_g E_{g,t}^{\frac{-1}{\sigma_e}} = p_{g,t} \tag{1.24}$$

where $p_{k,t}$ is the price paid for energy variety $k \in \{d, g\}$. Finally, for the labour market to clear, the price of the energy varieties must be such that

$$p_{k,t} = \frac{w_t}{A_{k,t}} \quad k \in \{d, g\} \tag{1.25}$$

As markets are competitive and all production functions are constant returns to scale, profits in the economy are zero.

The middle aged choose their household consumption and savings such that their Euler equation holds:

$$\frac{1}{c_t^{t-1}} = \left(1 + \theta_t^y \frac{1 + n_t}{\psi_t^{t-1}}\right) \psi_{t+1}^{t-1} \beta \mathbb{E}_t \left[\frac{R_{t+1}}{\psi_{t+1}^{t-1}} \frac{1}{c_{t+1}^{t-1}}\right] \tag{1.26}$$

This Euler equation has the standard form - the marginal utility of consumption when middle-aged is equated to the expected marginal utility of saving - but with some minor adjustments. First, since each additional unit of middle aged consumption increases total household consumption by $1+\theta_t^y\frac{1+n_t}{\psi_t^{t-1}}$, the return to saving on the right-hand side of the Euler equation needs to be multiplied by this factor. $1+\theta_t^y\frac{1+n_t}{\psi_t^{t-1}}$ is increasing in the ratio of child-to-parent consumption (θ_t^y) and the child-to-parent ratio ($\frac{1+n_t}{\psi_t^{t-1}}$). Additionally, since the middle-aged survive to retirement with probability ψ_{t+1}^{t-1}, the benefits of saving need to be weighted by this term. Finally, due to the presence of perfect annuity markets that redistribute the assets of those who die before retirement, the return on capital conditional on surviving to retirement is given by $\frac{R_{t+1}}{\psi_{t+1}^{t-1}}$.

1.3.2 Laissez-Faire Equilibrium

The laissez-faire equilibrium is given by the competitive equilibrium in which $\tau_t = d_t^{t-1} = d_t^{t-2}$=0 in all periods. As shown in Section A.1 of the Appendix, the labour allocations are solved as

$$L_{y,t} = \left(\frac{1-\nu-\alpha}{1-\alpha}\right) N_t^{t-1} \tag{1.27}$$

$$L_{k,t} = \left(\frac{\nu}{1-\alpha}\right)\left(\frac{\omega_k^{\sigma_e} A_{k,t}^{\sigma_e-1}}{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e-1}}\right) N_t^{t-1} \quad k \in \{d, g\} \tag{1.28}$$

The allocation of labour across the final goods sector and the energy sector as a whole is determined by the Cobb-Douglas parameters of the final goods production function. The allocation of labour within the energy sector is determined by the weights on clean and dirty energy (ω_k), and relative productivity ($\frac{A_{g,t}}{A_{d,t}}$) across the two energy types. Consequently, when $\sigma > 1$ (as I assume in the calibrated version of my model), the labour allocation grows (shrinks) in the energy sector with higher (lower) TGP growth.

The labour allocations then determine $E_{d,t}$, $E_{g,t}$ and $E_{y,t}$. With S_{t-1} given as a state variable and $E_{d,t}$ solved, (1.18) determines S_t. k_t^{t-2} is given as a state variable, which determines $K_{y,t}$. Aggregate output is then pinned down from (1.8). Combining the Euler equation (1.26) and budget constraints of the middle aged (1.6 and 1.7), the consumption and savings choices of the middle-aged households are given by

$$c_{h,t} = \frac{1}{1+\psi_{t+1}^{t-1}\beta} w_t \tag{1.29}$$

$$c_t^{t-1} = \left(1+\theta_t^y \frac{1+n_t}{\psi_t^{t-1}}\right)^{-1} c_{h,t} \tag{1.30}$$

$$c_t^t = \theta_t^y c_t^{t-1} \tag{1.31}$$

$$k_{t+1}^{t-1} = \frac{\psi_{t+1}^{t-1}\beta}{1+\psi_{t+1}^{t-1}\beta} w_t \tag{1.32}$$

while retiree consumption is

$$c_t^{t-2} = \frac{R_t}{\psi_t^{t-2}} k_t^{t-2} \tag{1.33}$$

The fact that middle-aged agents choose to save a fraction $\frac{\psi_{t+1}^{t-1}\beta}{1+\psi_{t+1}^{t-1}\beta}$ of their lifetime income for retirement follows from the assumptions that their household's consumption is proportional to their individual consumption and that they have log utility preferences over their individual consumption. As in a standard OLG model, the marginal propensity to save out of household income is increasing in β. An increase in the life expectancy of retirees (ψ_{t+1}^{t-1}) similarly increases the savings rate of middle-aged households by extending their expected time horizon. Factor prices (1.20)-(1.22) and the marginal propensity to consume of middle-aged households (1.29) imply an aggregate savings rate given by

$$s_t \equiv \frac{K_{y,t+1}}{Y_t} = (1-\alpha)\left(\frac{\psi_{t+1}^{t-1}\beta}{1+\psi_{t+1}^{t-1}\beta}\right) \tag{1.34}$$

The aggregate savings rate is the weighted average of the saving rates of middle-aged households (given by $\frac{\psi_{t+1}^{t-1}\beta}{1+\psi_{t+1}^{t-1}\beta}$) and retired households (who save nothing). The weight given to the savings rate of middle-aged households is $1-\alpha$, which is the labour share of income. An increase in either the middle-aged savings rate or their income share increases the aggregate savings rate.

1.3.3 Optimal Carbon Tax

This sub-section characterises the optimal carbon tax in a second best scenario in which the policymaker is unable to use the lump sum rebates d_t^s to alter the intratemporal distribution of income across agents. That is, in any given period, the distribution of income between the middle aged and retirees must remain unaltered. In the context of this model, this restriction is equivalent to the restriction that the distribution of income between capital and labour must remain unaltered by the carbon tax. The retiree income share is given by α in the laissez-faire equilibrium. With climate policy, their income share is given by

$$\frac{\left(k_t^{t-2}\frac{R_t}{\psi_t^{t-2}} + d_t^{t-2}\right)N_t^{t-2}}{Y_t} = \alpha + \frac{d_t^{t-2}N_t^{t-2}}{Y_t}$$

This share is equal to α iff $d_t^{t-2} = 0$ and the middle aged receive all revenues generated via the carbon tax. This assumption is necessary for the derivation of a closed-form solution of the optimal tax-to-GDP ratio, but is also not unreasonable given that carbon pricing schemes are often formulated without specific distributional aims in mind with regards to the distribution of income across generations or capital and labour at a given point in time. Despite this restriction, the policymaker is still able to affect the intergenerational distribution of welfare given the intertemporal impacts of climate policy. This means that the optimal carbon tax formula is second best: it partly reflects the desire to internalise the climate externality and it partly reflects a desire to redistribute income across generations.

I assume that the policymaker has the following utilitarian social welfare function

$$\mathbb{E}_t\left[U_t\right] = \mathbb{E}_t\left[\left(\frac{\beta}{\phi}\right)^2 N_t^{t-2}\ln c_t^{t-2} + \left(\frac{\beta}{\phi}\right)\left(N_t^{t-1}\ln c_t^{t-1} + N_{t+1}^{t-1}\beta\ln c_{t+1}^{t-1}\right) + \sum_{s=t}^{\infty}\phi^{s-t}\sum_{i=0}^{2}N_{s+i}^{s}\beta^i\ln c_{s+i}^{s}\right] \tag{1.35}$$

That is, expected social welfare at time t is the sum of the expected lifetime welfare of each generation, with lifetime utility across generations discounted using the pure discount factor[12] (ϕ). Thus, in addition to the pure intragenerational rate of time preference implied by $\beta \leq 1$ in individual preferences, there is a pure intergenerational rate of time preference implied by $\phi \leq 1$ in the policymaker's preferences. The objective of the policymaker is to set carbon taxes and lump sum transfers to maximise $\mathbb{E}_t\left[U_t\right]$ subject to the competitive equilibrium constraints given by equations (1.20)-(1.26).

Lemma 1: The policymaker's problem in primal form is to choose a sequence of allocations $\left\{c_s^{s-1}, k_{s+1}^{s-1}, E_{y,s}, K_{y,s}, L_{y,s}, E_{d,s}, E_{g,s}, L_{d,s}, L_{g,s}, S_s\right\}$ to maximise

$$\sum_{s=t}^{\infty}\phi^{s-t}\mathbb{E}_t\left[\tilde{N}_s\ln c_s^{s-1}\right]$$

[12]Weighting the utility of the current middle aged and retired generations by $\frac{\beta}{\phi}$ and $\left(\frac{\beta}{\phi}\right)^2$ has no substantive impact on the results but simplifies notation.

subject to the following constraints:

$$
\begin{aligned}
c_s^{s-1}\bar{N}_s + K_{y,s+1} &= e^{-\gamma_s\left(S_s-\bar{S}\right)} A_{y,s} E_{y,s}^{\nu} K_{y,s}^{\alpha} L_{y,s}^{1-\nu-\alpha} \\
c_s^{s-1}\bar{N}_s &= (1-s_s)\, e^{-\gamma_s\left(S_s-\bar{S}\right)} A_{y,s} E_{y,s}^{\nu} K_{y,s}^{\alpha} L_{y,s}^{1-\nu-\alpha} \\
E_{y,s} &= \left(\omega_d E_{d,s}^{\frac{\sigma_e-1}{\sigma_e}} + \omega_g E_{g,s}^{\frac{\sigma_e-1}{\sigma_e}}\right)^{\frac{\sigma_e-1}{\sigma_e}} \\
E_{k,s} &= A_{k,s} L_{k,s} \quad k \in \{d, g\} \\
N_s^{s-1} &= L_{y,s} + L_{d,s} + L_{g,s} \\
S_s - \bar{S} &= \sum_{i=0}^{s} (1-d_i)\, E_{d,s-i}
\end{aligned}
$$

where $\tilde{N}_s \equiv N_s^s + \left(\frac{\beta}{\phi}\right) N_s^{s-1} + \left(\frac{\beta}{\phi}\right)^2 N_s^{s-2}$ is an adjusted measure of population size.

Proof: See Section A.2 of the Appendix.

Proposition 1: The optimal carbon tax to GDP ratio is given by

$$
\frac{\tau_t^*}{Y_t^*} = \sum_{s=t}^{\infty} \mathbb{E}_t \left[(1-d_{s-t})\,\gamma_s \quad \phi^{s-t} \frac{\tilde{N}_s}{\tilde{N}_t} \frac{(1-s_t)}{(1-s_s)} \; \frac{\left\{1 - \frac{(1-s_s)\Theta_s}{\tilde{N}_s}\right\}}{\left\{1 - \frac{(1-s_t)\mathbb{E}_t[\Theta_t]}{\tilde{N}_t}\right\}} \right]
$$

where

$$
\Theta_s \equiv \sum_{i=s}^{\infty} (\alpha\phi)^{i-s} \left(\frac{s_i}{1-s_i} \tilde{N}_i - \frac{\phi\alpha}{1-s_{i+1}} \tilde{N}_{i+1} \right)
$$

Proof: See Section A.3 of the Appendix.

Since $Y_t^* \approx Y_t$, where Y_t is GDP with the existing (suboptimal) carbon tax policy, we get the nearly closed form solution[13] for τ_t^* in terms of observables and exogenous variables/parameters given by

$$
\tau_t^* \approx Y_t \sum_{s=t}^{\infty} \mathbb{E}_t \left[\underbrace{(1-d_{s-t})\,\gamma_s}_{\text{Output Damages}} \quad \underbrace{\phi^{s-t} \frac{\tilde{N}_s}{\tilde{N}_t} \frac{(1-s_t)}{(1-s_s)}}_{\text{MRS}} \; \underbrace{\frac{\left\{1 - \frac{(1-s_s)\Theta_s}{\tilde{N}_s}\right\}}{\left\{1 - \frac{(1-s_t)\mathbb{E}_t[\Theta_t]}{\tilde{N}_t}\right\}}}_{\text{Second-Best Adjustment}} \right] \tag{1.36}
$$

[13] I use this approximation formula for illustrative purposes. For computational purposes, I iteratively solve the model to find the exact solution. I do so by first solving for τ_t^* using formula (1.36). I then solve the time t equilibrium of the model conditional on this estimate of τ_t^* to obtain a new value of Y_t consistent with this policy. I then use this revised value of Y_t to update my estimate of τ_t^*. I iterate this process until the solution converges to a sufficient degree of precision. All optimal carbon tax computations in this chapter are computed in a similar manner.

The interpretation of this formula is intuitive. The optimal carbon tax is equal to the present discounted value of the externality cost of emissions with an adjustment for the fact that the resulting equilibrium is still second best. The percentage loss in output due to an additional unit of atmospheric carbon concentration in period s is given by γ_s. Because carbon released into the atmosphere depreciates over time, this term must be multiplied by $1 - d_{s-t}$, which is the fraction of period t emissions still in the atmosphere in period s. To value this output loss in present utility terms, these damages are discounted using the policymaker's intertemporal marginal rate of substitution. This marginal rate of substitution has three components. First, the ϕ^{s-t} term reflects the policymaker's pure rate of time preference across generations. Second, $\frac{\tilde{N}_s}{\tilde{N}_t}$ is the growth rate of $\tilde{N}_s$ between period t and s. The growth rate of $\tilde{N}_s$ is relevant since the social welfare function is an aggregate of individual utility functions that have decreasing marginal utility in consumption. Consequently, all else equal, output losses lead to a greater loss of social welfare when the economy is populated by a greater number of agents due to the higher marginal utility of consumption at the individual level. When $\frac{\tilde{N}_s}{\tilde{N}_t}$ is greater than one (that is, $\tilde{N}_s$ is growing), this pushes against the discounting of the future due to the policymaker's pure rate of time preference. When $\beta \neq \phi$, the different generations alive in a given period are weighted differently by the policymaker, meaning that $\tilde{N}_s \equiv N_s^s + \left(\frac{\beta}{\phi}\right) N_s^{s-1} + \left(\frac{\beta}{\phi}\right)^2 N_s^{s-2}$ is the relevant measure of population size[14]. Third, the marginal rate of substitution is determined by $\frac{(1-s_t)}{(1-s_s)}$. This reflects the fact that conditional on output, the marginal utility of consumption is positively correlated with the aggregate savings rate. Periods with higher aggregate savings rates are thus periods in which consumption is relatively more valuable, meaning climate-related damages in these periods are given more weight. Finally, the expression $1 - \frac{(1-s_t)\mathbb{E}_t[\Theta_t]}{\tilde{N}_t}$ shows up in the formula because the rate of capital accumulation in the economy is socially inefficient according to the policymaker's social welfare function. The term Θ_s is equal to the discounted value of all current and future differences between the marginal product of capital and the policymaker's intertemporal marginal rate of substitution for aggregate consumption. The optimal carbon tax responds to changes in Θ_s over time as the policymaker can use the carbon tax to redistribute output to periods where consumption is undervalued in the competitive equilibrium. If the degree of this suboptimality does not change over time (that is, Θ_s is constant), there is no such rationale to adjust the carbon tax as consumption is equally scarce (or abundant) relative the first-best scenario in all periods.

Aggregate output in period s (Y_s) does not show up in the optimal tax formula because the output losses due to climate change are proportional to Y_s while the marginal utility of consumption and second best adjustments are inversely proportional to Y_s such that this

[14]In the special case when $\beta = \phi$, $\tilde{N}_s = N_s^s + N_s^{s-1} + N_s^{s-2} = N_s$.

term cancels out. Y_t appears in the formula because the marginal utility of consumption and second best adjustments in period t are proportional to Y_t. When Y_t is higher, the marginal utility of consumption in period t is lower, meaning that the policymaker is relatively more willing to substitute present consumption for future consumption. This is achieved by setting a higher carbon tax such that present emissions and output are reduced to limit climate change and increase future output.

Under the conditions that (i) $\phi = \beta$, (ii) the demographic parameters are all constant and population growth is equal to zero ($\psi_{t+1}^t = \psi_1$, $\psi_{t+2}^t = \psi_2$, $n_t = 0$), and (iii) the policymaker has an unrestricted ability to make lump sum transfers between agents and tax capital holdings, the optimal tax formula reduces to the Golosov et al. (2014) formula as a special case:

$$\tau_t^* \approx Y_t \sum_{s=t}^{\infty} \mathbb{E}_t \left[\underbrace{(1 - d_{s-t}) \, \gamma_s}_{\text{Output Damages}} \; \underbrace{\beta^{s-t}}_{\text{MRS}} \right] \tag{1.37}$$

1.4 Quantitative Analysis

This section discusses the results from the calibrated version of the model. In addition to quantifying the optimal carbon tax and its evolution over time, I illustrate the implications for climate change and welfare under the various policies. All results are derived under the assumption of perfect foresight. To illustrate the potential impact of uncertainty about future demographic change, I run all of my results using three different forecasts of the population by age group provided by the United Nations (2019).

1.4.1 Calibration

One period in the model has a duration of thirty years. As a general note, where data exists for a stock variable, it is mapped onto model periods by identifying the stock value at the mid-point of the period (for example, the population level during the 1950-1980 period in the model is set to the population level in 1965); where data exists for a flow variable, the flow is aggregated over the thirty years of the model period to obtain the corresponding flow value for that period (for example, the GDP level during the 1950-1980 period in the model is set to the sum of GDP from 1950 to 1979).

Both the pure discount factors (β and ϕ) are set equal to 0.985^{30} in the baseline calibration (see Section 1.5.1 for the implications of varying these two parameters independently). The Cobb-Douglas production function shares are set to achieve an energy share in GDP of 5%

and a 60%-40% pre-tax income split between labour and capital respectively. The production function parameters of the energy composite are set based on the recent study by Papageorgiou et al. (2017) and the values used in Golosov et al. (2014). Based on the average value of the various estimates and assumptions across two papers, the weight on dirty energy (ω_d) is set to 0.5 while the elasticity of substitution between dirty and green energy (σ_e) is set to 1.3.

The parameters of the carbon cycle (φ_L, φ_0, φ) and damage function (γ, $\bar{S}$) are set in line with those used in Golosov et al. (2014). Since Golosov et al. (2014) use a period length of 10 years while I use a period length of 30 years, the carbon cycle parameters φ_0 and φ need to be adjusted accordingly. Letting objects with overlines denote the Golosov et al. (2014) equivalents, I set φ_0 and φ such that

$$1 - d_i = 1 - \bar{d}_{1+3i}$$

which implies

$$\begin{aligned} \varphi_0 &= \bar{\varphi}_0 \left(1 - \bar{\varphi}\right) \\ (1 - \varphi) &= \left(1 - \bar{\varphi}\right)^3 \end{aligned}$$

All demographic data is obtained from the United Nations (2019) World Population Prospects. The initial population levels by age group in the 1950-1980 period are obtained by aggregating the United Nations (2019) World Population Prospects data by age into the relevant age categories, with all ages above 60 including those in the 90+ category aggregated into the 60-90 age group in the model. The survival probabilities of the young and middle aged in a given period (ψ_{t+1}^{t} and ψ_{t+1}^{t-1}) are then imputed from the time series data for life expectancies at ages 30 and 60, respectively. The UN provides projections for these life expectancies up to 2100. These projections are extended by assuming that life expectancy at 30 and 60 increase at a fixed growth rate to 60 and 30 respectively by 2200, from which point they remain constant. Letting $LE_{30,t}$ refer to life expectancy at 30 at time t and $LE_{60,t}$ be life expectancy at 60 at time t, I then set ψ_{t+1}^{t} and ψ_{t+1}^{t-1} such that

$$\begin{aligned} \psi_{t+1}^{t-1} &= \frac{LE_{60,t}}{30} \\ \psi_{t+1}^{t} &= \frac{LE_{30,t}}{30 + LE_{60,t+1}} \end{aligned}$$

Figure 1.2 illustrates the implied life expectancy paths. Unlike the other population data I use, the UN only provides the medium variant projections for life expectancy at ages 30 and 60. The results from Section 1.4.2 suggest that the optimal policy is robust to changes in ψ^t_{t+1} and ψ^t_{t+2} conditional on n_t, so only having one set of life expectancy projections to work with does not seem to be major a constraint.

Figure 1.2: Global Life Expectancy at Ages 30 and 60 (Years)

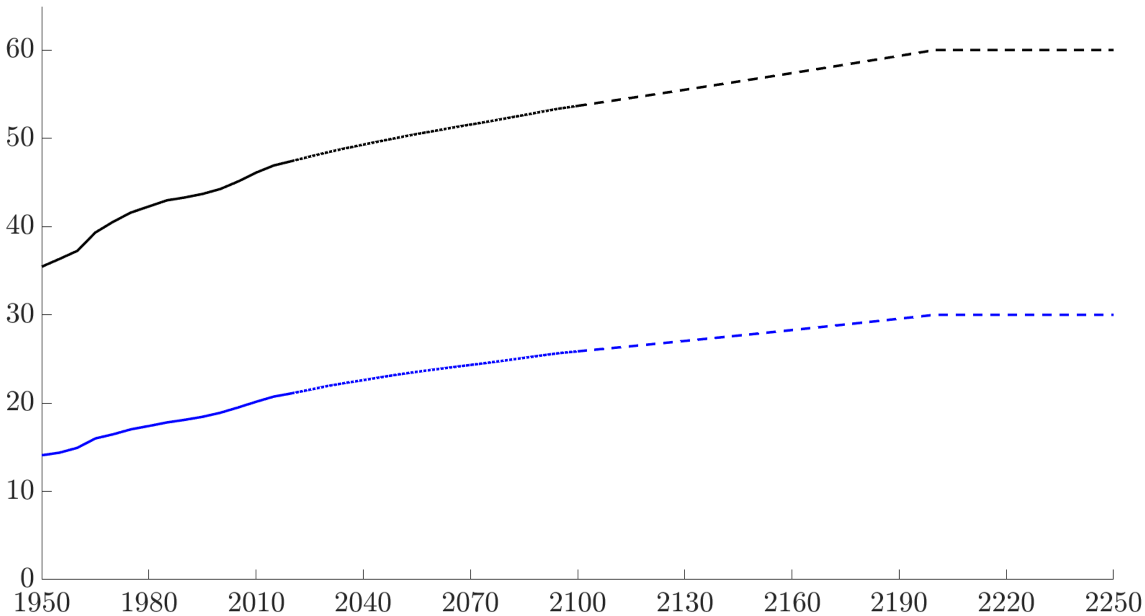

Notes: The blue and black lines refer to life expectancy at ages 60 and 30, respectively. The solid portions of the lines (1950-2020) represent historical estimates; the dotted portions (2020-2100) are the UN Medium variant projections; and the dashed portions (2100-2250) are my own extrapolations of the UN projections. Source: United Nations (2019); Author's calculations.

The birth rate n_t is then chosen to ensure that the total population size implied by the initial population levels by age and the chosen values for n_t, ψ^t_{t+1} and ψ^{t-1}_{t+1} produces a time series that matches the population size estimates and projections in the UN data. As with the life expectancy data, the population data projections provided by the UN only run to 2100. To extend these projections, I assume that the projected gross growth rate of the population in 2100 changes at a constant rate such that it equals 1 in 2200, after which I assume the population remains constant. As the UN provides various projections for the total population size, I do this exercise for the low, medium and high projection variants and run all model results using these three extended projections. Figure 1.3 illustrates these extended projections.

The ratio of young to middle-aged consumption (θ^y_t) is set equal to 1. Provided this ratio is exogenously fixed, its level does not affect any of the results discussed in this chapter given the other assumptions I make. In the absence of strong evidence to the contrary, I therefore

Figure 1.3: Global Population (Billions)

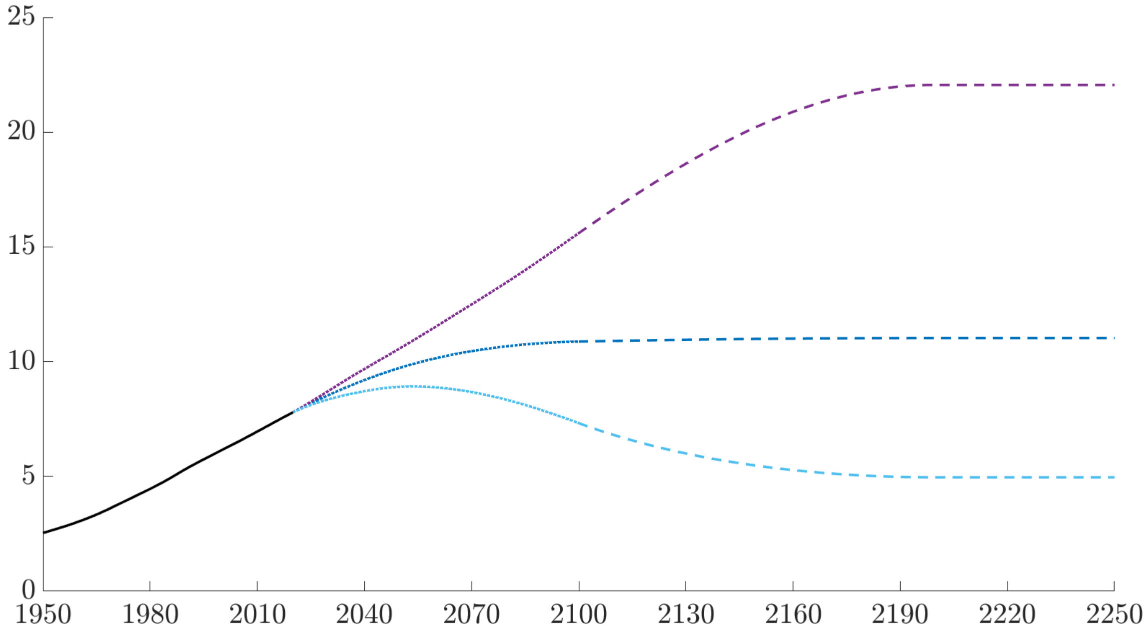

Notes: The solid portions of the lines (1950-2020) represent historical estimates; the dotted portions (2020-2100) are the various UN projections; and the dashed portions (2100-2250) are my own extrapolations of these UN projections. The light blue, dark blue, and purple dashed lines reflect the Low, Medium and High projection variants, respectively.
Source: United Nations (2019); Author's calculations.

choose a value of unity as an intuitively reasonable benchmark.

To calibrate S_{t-1}, I follow Golosov et al. (2014) in expressing S_t as the sum of two components: a permanent component (denoted by $S_{1,t}$) and a transitory component (denoted by $S_{2,t}$). It follows from (1.18) and (1.19) and these definitions of $S_{1,t}$ and $S_{2,t}$ that we can write these components recursively as

$$S_{1,t} = S_{1,t-1} + \varphi_L E_{d,t} \tag{1.38}$$

$$S_{2,t} = (1-\varphi) S_{2,t-1} + (1-\varphi_L) \varphi_0 E_{d,t} \tag{1.39}$$

Using this recursive representation, I calculate the initial condition on S_{t-1} by assuming that all atmospheric carbon in the pre-industrial era is permanent ($\bar{S}_1 = \bar{S}$ and $\bar{S}_2 = 0$) and then running historical data on global CO_2 emissions from fossil fuels from the Global Carbon Project (2021) through equations (1.38) and (1.39) to arrive at an estimate for S_{t-1} in the relevant period[15].

[15]Note that the data is measured in terms of CO_2 emissions. Since emissions and carbon concentrations in the model are expressed in terms of carbon, CO_2 emissions are converted into carbon equivalents by using the standard conversion factor of $\frac{12}{44}$. That is, one ton of CO_2 is assumed to contain $\frac{12}{44}$ tons of carbon. These

Finally, the TFP series ($A_{y,t}$, $A_{d,t}$, $A_{g,t}$) are calibrated as follows. For the first two periods of the model (1950-1980 and 1980-2010), these three variables are pinned down by matching the model output to the data on real GDP, aggregate emissions, and the relative price between dirty and green energy assuming that no carbon taxes are active in the model. This is done by starting with a target relative price between dirty and green energy. Equation (1.25) for dirty and green energy implies that

$$\frac{\bar{p}_{g,t}}{\bar{p}_{d,t}} = \frac{A_{d,t}}{A_{g,t}} \tag{1.40}$$

and thus pins down the relative TFP across the two types of energy. Plugging (1.40) into (1.28) for the dirty energy sector implies

$$L_{d,t} = \left(\frac{\nu}{1-\alpha}\right)\left(\frac{\omega_d^{\sigma_e}}{\omega_d^{\sigma_e} + \omega_g^{\sigma_e}\left(\frac{\bar{p}_{d,t}}{\bar{p}_{g,t}}\right)^{\sigma_e - 1}}\right) N_t^{t-1} \tag{1.41}$$

Given a target value for dirty energy emissions $\bar{E}_{d,t}$, taken from the Global Carbon Project (2021) databased on fossil fuel emissions, $A_{d,t}$ is then pinned down by plugging (1.41) into (1.11)

$$A_{d,t} = \left[\left(\frac{\nu}{1-\alpha}\right)\left(\frac{\omega_d^{\sigma_e}}{\omega_d^{\sigma_e} + \omega_g^{\sigma_e}\left(\frac{\bar{p}_{d,t}}{\bar{p}_{g,t}}\right)^{\sigma_e - 1}}\right) N_t^{t-1}\right]^{-1} \bar{E}_{d,t}$$

Finally, I pin down $A_{y,t}$ using global estimates of real GDP in 2015 US Dollar terms provided by the World Bank (denoted here by $\bar{Y}_t$). Conditional on $\bar{Y}_t$, $A_{y,t}$ is then pinned down from (1.8) and (1.9) as

$$A_{y,t} = \left[e^{-\gamma_t\left(S_t - \bar{S}\right)} E_{y,t}^{\nu} K_{y,t}^{\alpha} L_{y,t}^{1-\nu-\alpha}\right]^{-1} \bar{Y}_t$$

The endogenous terms on the right hand side of this equation are calculated as follows. The initial $K_{y,t=1}$ for the period 1950-1980 is estimated as

$$K_{y,t=1} = s_{t=1}\bar{Y}_{t=0} \approx s_{t=0}\bar{Y}_{t=0}$$

where s_t is calculated from (1.34). Subsequent values of $K_{y,t}$ are then determined as usual by $K_{y,t} = s_t Y_t$. $L_{y,t}$ is determined by (1.27). $E_{y,t}$ is pinned down first by combining (1.40), (1.28) and (1.11) to determine $E_{g,t}$, and then combining this with the $\bar{E}_{d,t}$ target in (1.10).

For all subsequent periods of the model (that is, from the 2010-2040 period onwards), I

adjustments are used both in setting S_{t-1} and also the emissions targets $\bar{E}_{d,t}$, discussed below. For more details on conversion factors, see United States Environmental Protection Agency (2022).

assume that $A_{d,t}$ remains constant while $A_{g,t}$ grows at a constant annual rate of 2%. Finally, I calibrate $A_{y,t}$ such that real GDP per capita would grow at an annual rate of 2% holding fixed climate-related damages at their 1980-2010 level.

The details of the calibration are summarised below in Table 1.1

Table 1.1: Model Calibration

Parameter/Variable	Description	Value	Target/Source
β	Intragenerational discount factor	0.985^{30}	Standard
ϕ	Intergenerational discount factor	0.985^{30}	Standard
$A_{y,t}$	Aggregate TFP		Real GDP
α	Cobb-Douglas share on capital	0.40	Capital share
ν	Cobb-Douglas share on energy	0.05	Energy share
$A_{d,t}$, $A_{g,t}$	Dirty and Green Energy TFP		$\frac{p_{d,t}}{p_{g,t}}$, Carbon emissions
ω_d	Energy weight on dirty energy	0.50	Papageorgiou et al. (2017)
σ_e	Energy elasticity of substitution	1.3	Papageorgiou et al. (2017)
γ	Elasticity of damage function	5.30×10^{-5}	Golosov et al. (2014)
$\bar{S}$	Pre-industrial GtC concentration	581	Golosov et al. (2014)
φ_L, φ_0, φ	Carbon cycle parameters	0.20, 0.38, 0.07	Golosov et al. (2014)
θ_t^y	Young consumption ratio	1	No prior
N_1^1, N_1^0, N_1^{-1}	Initial population levels		Population size
ψ_{t+1}^t, ψ_{t+1}^{t-1}	Survival probabilities		Life expectancy
n_t	Birth rate		Population size

Notes: Model periods are thirty years.

1.4.2 Optimal Carbon Tax

Figure 1.4, below, illustrates the extent to which the optimal carbon tax implied by the model deviates from the representative agent benchmark. The solid lines show the percentage deviation between the optimal carbon tax implied by (1.36) and the representative agent benchmark in (1.37) for the various UN population projection variants. The dashed lines calculate the percentage difference with the representative agent benchmark using the following equation:

$$\tau_t^* \approx Y_t \sum_{s=t}^{\infty} \mathbb{E}_t \left[\underbrace{(1 - d_{s-t})\, \gamma_s}_{Output\ \ Damages} \quad \underbrace{\phi^{s-t} \frac{\tilde{N}_s}{\tilde{N}_t}}_{MRS} \right] \tag{1.42}$$

The optimal carbon tax expression in (1.42) shuts down all demographic effects on the optimal carbon tax other than the one operating directly through fluctuations in $\tilde{N}_s$. The results show that accounting for demographic change can have significant impacts on the optimal carbon tax and that this impact is primarily accounted for by fluctuations in $\tilde{N}_s$.

The optimal carbon tax in the 2010-2040 period of the model using the representative agent formula (1.37) is \$61 per ton of CO_2 (or \$222 per ton of carbon). Using the full optimal tax formula from this chapter (1.36), the optimal carbon tax is 9.3% lower than the representative agent benchmark under the low variant population projection, and 9.5% and 36% higher using the medium and high variant population projections, respectively. These gaps all converge to zero over time as the model reaches a demographic steady state and starts behaving as though the economy is populated by a representative agent. These results are intuitive. Because the marginal utility of each agent is decreasing in their consumption, a given level of aggregate consumption leads to higher social welfare when it is spread across a larger population. When the global population is expected to grow, this means that more weight should be placed on the future than the present, all else equal. This effect partially offsets the discounting of the future due to pure rates of time preference and consumption growth present in representative agent IAMs. Since most benefits of carbon taxation accrue in the future while the costs are borne in the present, discounting the future less heavily implies that a higher carbon tax should be set at any given point in time. As the demographic shift slows, however, and the population size and structure stabilises, the optimal tax converges to the one implied by a similarly calibrated representative agent model. This convergence means that the optimal carbon tax with overlapping generations must grow less quickly than under the equivalent representative agent framework. It is worth noting here that with constant, positive population growth ($n_t = n > 0$), constant life expectancy ($\psi^t_{t+1} = \psi_1$, $\psi^t_{t+2} = \psi_2$), and $\beta = \phi$, the optimal carbon tax is permanently n% higher than the one implied by the representative agent formula. The declining wedge over time between the two tax formulae in Figure 1.4 is therefore driven by the fact that the projected population growth rates are initially non-zero and that they converge to zero.

The fact that the discrepancies between my OLG model and an equivalent representative agent framework are driven by the effect of population growth on the policymaker's discount rate is robust to the calibration of the model conditional on the functional forms chosen for the utility and damage functions. Because marginal damages from emissions are proportional to aggregate output while agents' marginal utility is inversely proportional to aggregate output, discounted future damages are not a function of future aggregate output. This means that demographic changes cannot affect the optimal carbon tax via aggregate output changes. In addition, the impact of demographic changes on the optimal carbon tax via savings rate fluctuations tends to be small. Increases in life expectancy over time increase the savings rate (see Figure 1.5), pushing up future marginal utility for a given level of output. This implies that the policymaker would like to set a higher carbon tax now, as future output damages are more painful given the increased need to save induced by the

Figure 1.4: Difference in Optimal Tax from Representative Agent Benchmark (%)

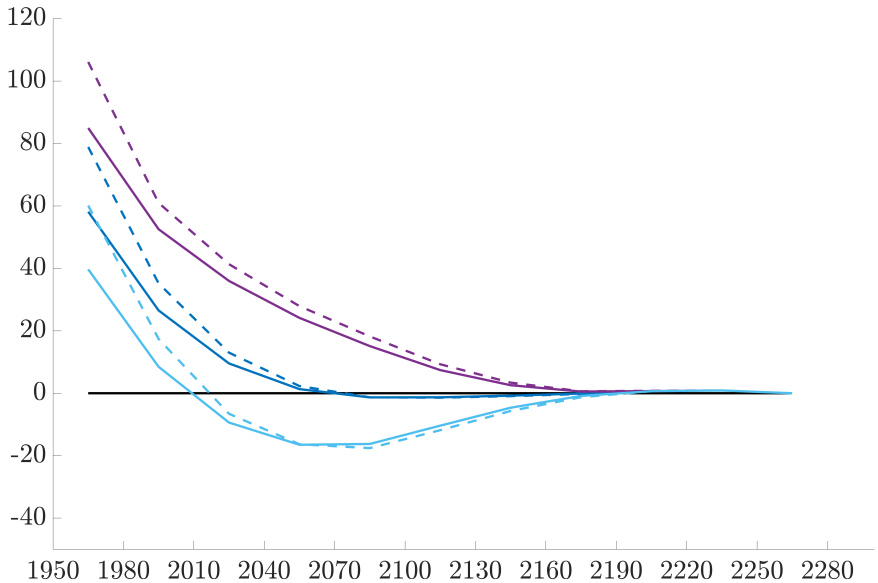

Notes: The solid lines represent the percentage deviation between the optimal carbon tax implied by the OLG formula (1.36) and the equivalent representative agent formula (1.37); the dashed lines represent the percentage deviation between the OLG formula that accounts only for fluctuations in $\tilde{N}_s$ (1.42) and the equivalent representative agent formula (1.37). The light blue, dark blue, and purple lines reflect the results of calibrating the model to the low, medium and high United Nations (2019) population projection variants, respectively.

Table 1.2: Percentage Difference in Optimal Tax from Representative Agent Benchmark

	L Variant Projection				M Variant Projection				H Variant Projection			
	(i)	(ii)	(iii)	(iv)	(i)	(ii)	(iii)	(iv)	(i)	(ii)	(iii)	(iv)
1950-1980	39.7	60.1	-0.7	-15.0	58.1	78.9	-0.7	-13.7	84.9	106.1	-0.7	-12.3
1980-2010	8.5	17.4	-1.0	-11.9	26.5	35.2	-1.0	-10.5	52.5	61.0	-1.0	-8.9
2010-2040	-9.4	-6.6	-1.5	-7.5	9.5	13.0	-1.5	-7.3	36.0	41.3	-1.5	-7.3
2040-2070	-16.5	-16.3	-1.6	-4.0	1.3	2.3	-1.6	-4.9	24.1	27.9	-1.6	-6.2
2070-2100	-16.2	-17.6	-1.5	-1.4	-1.3	-1.3	-1.5	-3.4	15.1	18.2	-1.5	-5.4
2100-2130	-10.4	-11.8	-1.1	-0.7	-1.3	-1.4	-1.1	-2.4	7.4	9.3	-1.1	-4.0
2130-2160	-4.6	-5.7	-0.6	-0.5	-0.8	-0.9	-0.6	-1.5	2.6	3.4	-0.6	-2.4
2160-2190	-0.9	-1.3	0.0	-0.1	-0.1	-0.2	0.0	-0.4	0.5	0.6	0.0	-0.6
2190-2220	0.6	0.7	0.5	0.5	0.7	0.8	0.5	0.4	0.7	0.8	0.5	0.4
2220-2250	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9
2250-2280	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Notes: The L, M and H Variant Projection columns refer to calculations done using the extended Low, Medium and High population size projections from the United Nations (2019). For each projection, column (i) refers to the results using the full optimal carbon tax formula; column (ii) refers to results when only adjusting for $\tilde{N}_t$ fluctuations; column (iii) refers to results when only adjusting for s_t fluctuations; column (iv) refers to results when only making second-best adjustments.

Figure 1.5: Projected Trends in Middle-Aged Survival Probability and Savings Rates

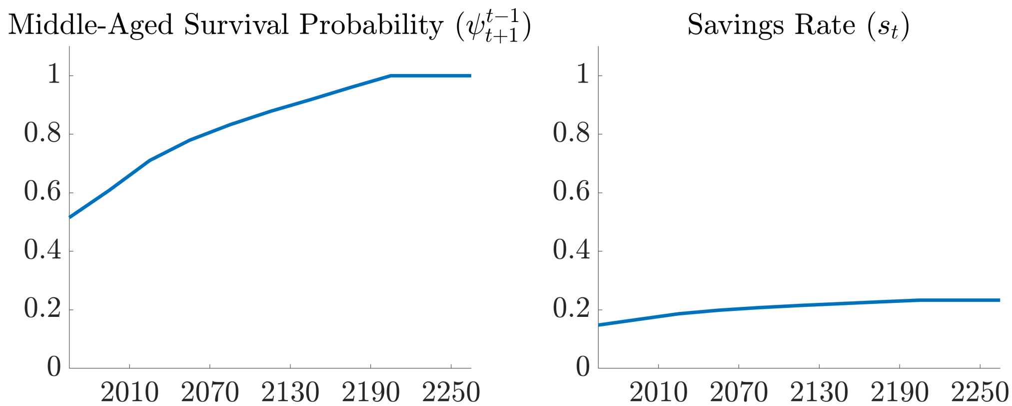

Figure 1.6: $\frac{(1-s_t)}{(1-s_s)}$ as a Function of ψ_{s+1}^{s-1}

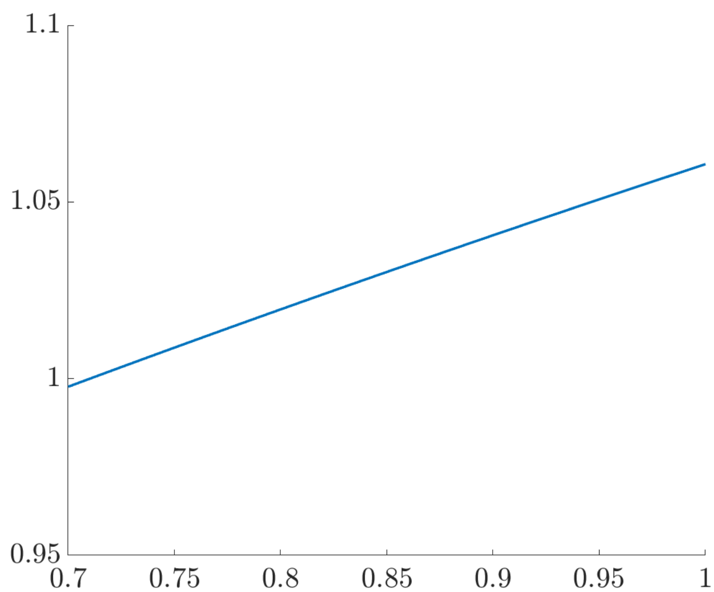

higher dependency ratio. The quantitative impact of this savings rate effect on the optimal carbon tax is small, however, for any reasonable calibration of the model. It is the ratio of the current to future marginal propensity to consume ($\frac{1-s_t}{1-s_s}$) that matters for how future climate damages are discounted. As illustrated in Figure 1.6, the projected increases in ψ_{s+1}^{s-1} over time do very little to alter this ratio. Moreover, the largest changes in this ratio occur furthest into the future and are thus discounted most heavily by the pure intergenerational discount factor ϕ.

In contrast, the optimal carbon tax is proportional to $\frac{\tilde{N}_s}{\tilde{N}_t}$ (rather than some non-linear function of this ratio), which is furthest from 1 in the near future before converging to 1 over time. This means that the channel that operates directly via the adjusted population growth rate dominates the one that operates through changes in life expectancy and savings rates.

Columns (ii)-(iv) in Table 1.2 provide a precise breakdown of how each channel contributes (in percentage point terms) to the wedge between the optimal carbon tax and the representative agent carbon tax. The contribution of all interactions effects between the three channels is given by the difference between column (i) and the sum of columns (ii)-(iv). As shown by this table, the effects operating directly through the growth rate of $\tilde{N}_t$, captured by column (ii), are the primary contributor to the overall wedge between the two tax rates in column (i).

1.4.3 Implications for Climate Change and Welfare

Figure 1.7 shows the evolution of carbon emissions ($E_{d,t}$) and atmospheric carbon concentration (S_t) under various population projection and policy scenarios. For each population projection, the figure illustrates the path of $E_{d,t}$ and S_t with no policy, the suboptimal carbon tax implied by the representative agent formula (1.37), and the optimal carbon tax (1.36). Absent policy, emissions continue to rise rapidly for a number of decades before naturally peaking and then declining. This increase in emissions is larger and lasts longer for population projections that forecast higher rates of population growth[16]. This decades-long increase in emissions causes a significant increase in atmospheric carbon concentrations, with S_t increasing to between 1,327 GtC and 2,175 GtC. In contrast, the introduction of a carbon tax leads to a sharp and immediate reduction in emissions, with S_t stabilising at around 1,000 GtC in all population projection scenarios.

Figure 1.8 shows how these changes in $E_{d,t}$ and S_t feed into temperature changes relative

[16] Cumulative emissions under all simulated scenarios remain well below the 5,000 GtC content of global fossil fuel reserves estimated by Rogner (1997). As fossil fuel resources are not exhausted, they are sold at the unit cost and do not earn a scarcity rent in the calibrated model.

Figure 1.7: Path of Emissions and Atmospheric Carbon Concentration

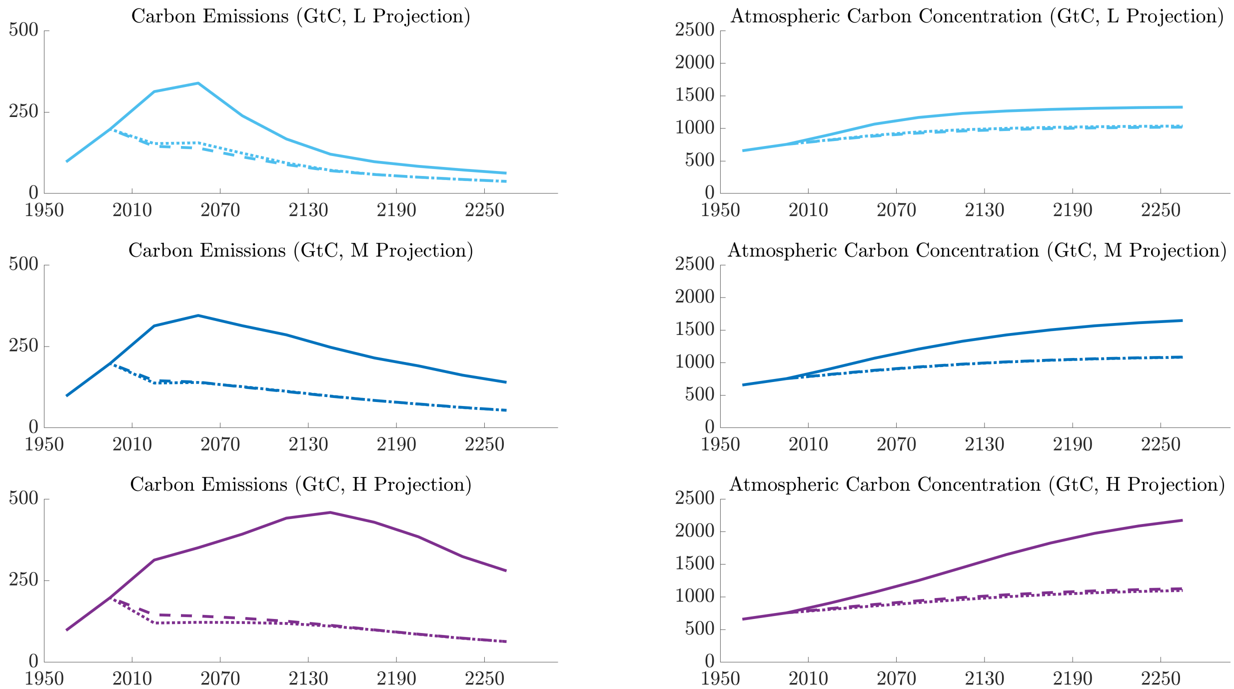

Notes: These figures illustrate how carbon emissions and the atmospheric carbon concentration evolve when there is (a) no carbon tax (b) the carbon tax that is optimal in the representative agent setting and (c) the optimal OLG carbon tax. Within each figure, the solid line represents scenario (a), the dashed line scenario (b), and the dotted line scenario (c). The three rows present the results using the low, medium and high variant United Nations (2019) global population projections. All values are measured in Gigatons of Carbon (GtC).

to pre-industrial levels (T_t) and climate-related damages (D_t). I follow Golosov et al. (2014) in mapping S_t into T_t using the following formula

$$T_t = 3\frac{\ln\left(\frac{S_t}{\bar{S}}\right)}{\ln 2}$$

This formula is a simple representation of the more complex relationships captured in climate models and provides a good approximation of the relationship observed in the data. As shown by Figure 1.8, the average global temperature change relative to pre-industrial levels is large absent policy, ranging from 3.6 degrees Celsius in the low variant projection to 5.7 degrees Celsius under the high variant projection. With carbon taxes in place, this temperature increase is moderated, and always lies in the range of 2.4 to 2.9 degrees Celsius. It is worth noting, however, that optimal warming under all population projections exceeds the 2 degree threshold for warming put in place by the Paris Agreement. In terms of climate-related damages, these are forecasted to be substantial absent climate policy, rising to as much as 8.1% of global GDP at the end of the simulation period under the high variant projection. With carbon taxes in place, damages are limited to the 2 to 3% range.

Figure 1.8: Path of Temperature Change and Climate Damages

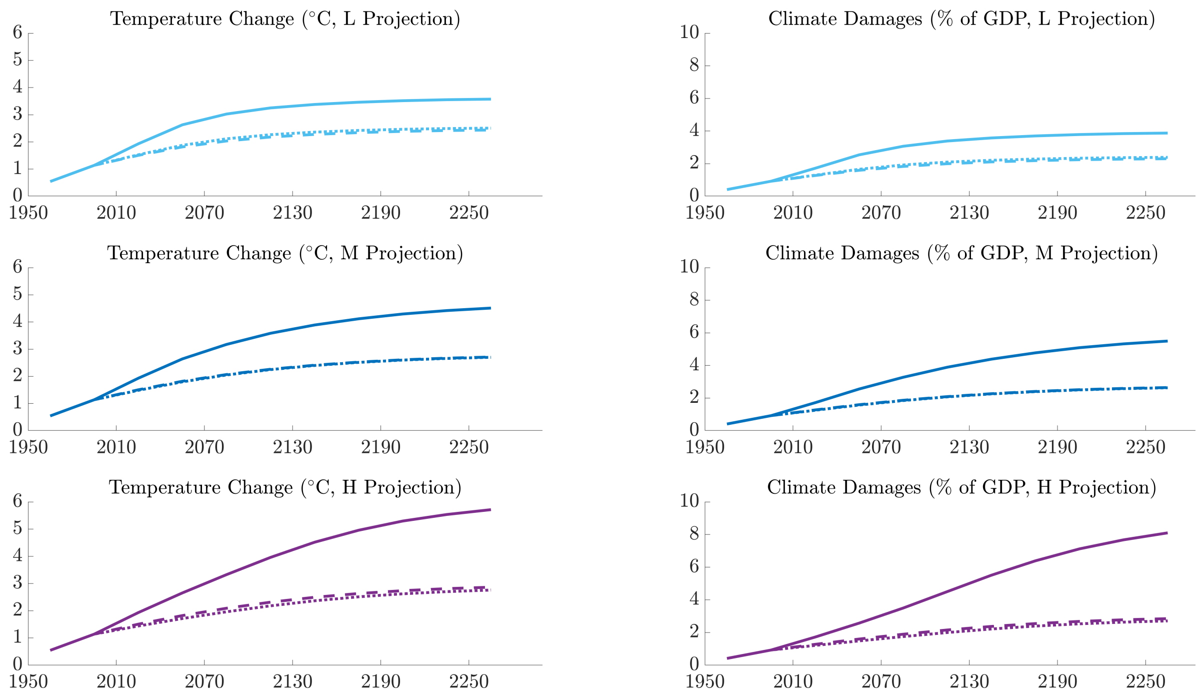

Notes: These figures illustrate how the average global temperature change and climate damages evolve when there is (a) no carbon tax (b) the carbon tax that is optimal in the representative agent setting and (c) the optimal OLG carbon tax. Within each figure, the solid line represents scenario (a), the dashed line scenario (b), and the dotted line scenario (c). The three rows present the results using the low, medium and high variant United Nations (2019) global population projections. Temperature changes are measured in the increase in average temperatures relative to the pre-industrial baseline, measured in degrees Celsius. Damages represent the percentage reduction in global GDP relative to a world where there is no temperature change.

Using the medium and high variants of the population projections, the optimal carbon tax implied by (1.36) is higher than that suboptimal tax implied by the representative agent formula (1.37). As a result, the optimal carbon tax leads to slightly lower emissions, temperature increases and damages than the suboptimal tax. This difference is more pronounced in the high variant simulation as the gap between the carbon taxes is larger. Conversely, under the low variant population projection, the optimal carbon tax is lower than the one implied by the representative agent formula, leading to higher emissions, temperature increases and damages under the optimal tax.

I calculate the welfare gains from the optimal tax by backing out how much additional aggregate consumption is required under a given suboptimal scenario to obtain a level of social welfare equal to the one obtained under the optimal policy (which I term \$2015 lump-sum aggregate consumption equivalent). Letting variables with stars denote outcomes under the optimal policy and variables without starts denote suboptimal outcomes, I first identify the required change in consumption of the middle aged in period t ($\Delta_{c,t}$) such that

$$U_t^* = U_t\left(\Delta_{c,t}\right)$$

Using the results from Appendix A.2, we can write

$$\Delta_{c,t} = c_t^{t-1}\left[\exp\left(\frac{U_t^* - U_t}{\tilde{N}_t}\right) - 1\right]$$

Letting C_t denote aggregate consumption, it thus follows that the change in aggregate consumption required to equate U_t^* and U_t (denoted by $\Delta_{C,t}$), is given by

$$\Delta_{C,t} = C_t\left[\exp\left(\frac{U_t^* - U_t}{\tilde{N}_t}\right) - 1\right]$$

The welfare gains associated with (i) moving from no carbon tax to the optimal carbon tax (denoted by $\Delta^1_{C,t}$) and (ii) moving from the suboptimal carbon tax to the optimal carbon tax (denoted by $\Delta^2_{C,t}$) are given in Table 1.3.

Table 1.3: Welfare Gains from Carbon Taxation ($2015, Trillions)

	L Variant Projection	M Variant Projection	H Variant Projection
$\Delta^1_{C,t}$	53.38	102.94	231.05
$\Delta^2_{C,t}$	0.63	0.15	2.17
$\frac{\Delta^2_{C,t}}{\Delta^1_{C,t}}$ (%)	0.94	0.14	1.17

Notes: The L, M and H Variant Projection columns refer to calculations done using the extended low, medium and high population size projections from the United Nations (2019). $\Delta^1_{C,t}$ is the welfare gain from implementing the optimal carbon tax relative to no tax; $\Delta^2_{C,t}$ is the welfare gain from implementing the optimal carbon tax relative to the suboptimal tax; the final row shows $\Delta^2_{C,t}$ as a percentage of $\Delta^1_{C,t}$.

The results show that while most welfare gains of the carbon tax can be achieved by using the suboptimal representative agent carbon tax, there are still significant welfare gains associated with moving from the suboptimal tax to the optimal tax ($148 Billion to $2.17 Trillion). $\Delta^2_{C,t}$ is largest under the high variant population projection for two reasons. First, under both the low and high variant projections, the optimal carbon tax differs more from the suboptimal tax than with the medium variant projection due to the larger and more persistent projected changes in the population size. Consequently, the representative agent tax formula is less precise and thus introduces larger policy mistakes. This explains why $\Delta^2_{C,t}$

is smallest for the medium variant projection. In addition, as shown in Figures 1.7 and 1.8, climate change is a larger issue the larger the projected increase in population size, making policy instrument errors of a given size more costly. Since these two factors both push in the same direction for the high variant projection, the welfare costs of the suboptimal carbon tax are largest in this case.

1.5 Discounting and Optimal Carbon Taxation

This section explores two key debates related to discounting and optimal carbon taxation. First, I revisit the longstanding Nordhaus-Stern debate regarding the appropriate pure rate of time preference to be used in optimal carbon tax calculations. Second, I assess the discounting debate in light of the more recent low interest rate environment.

1.5.1 Intra- Versus Inter-Generational Time Preference

One of the most notable sources of disagreement in the climate policy and optimal carbon taxation literature has been about the appropriate choice of the pure rate of time preference. Broadly speaking, two views have been put forward on this subject, often in opposition to one another. On one hand, there are positivist arguments, most famously championed by Nordhaus (2008), that argue for a pure rate of time preference informed by observed asset returns and consumption growth. In contrast, one can take a normative approach to discounting, as done by Stern (2007). According to this position, we have little ethical reason to treat generations differently, implying that we should use a much lower pure rate of time preference that is closer to zero. It is no surprise that the optimal carbon tax is very sensitive to the assumed pure rate of time preference given that a unit of emissions today produces output losses far into the future. The aim of this section of the chapter is not to take a particular stance in terms of what the "correct" pure rate of time preference should be; rather, I aim to show how this long-standing debate on discounting can be mapped to my model to illustrate some of the insights gained by moving away from the representative agent framework.

The setup of my model implies that there are two distinct pure time preference parameters in the policymaker's objective function: the intragenerational discount factor (β) and the intergenerational discount factor (ϕ). By introducing two pure time preference parameters, my model is able to simultaneously accommodate both the Nordhaus and Stern world views. Given that agents do not have bequest motives when saving for retirement, observables such as asset returns and consumption growth within my framework are informative about the pure intragenerational discount factor (β) but tell us nothing about how policymakers should value welfare across generations. The pure intergenerational discount factor (ϕ) is, therefore, a parameter that is free of positive implications (in the laissez-faire equilibrium, at least) and can thus be thought of purely in normative terms. Given this separation, I can test the robustness of the optimal policy to assumptions on each of these two parameters. Figure 1.9, below, illustrates the percentage change in the optimal carbon tax when ϕ (β) is raised from a minimum annualised value of 0.9 to a maximum annualised value of 0.999 for all values of β (ϕ) within the corresponding range $[0.9, 0.999]$.

Figure 1.9: Sensitivity of Optimal Tax to Time Preference (Percentage Change)

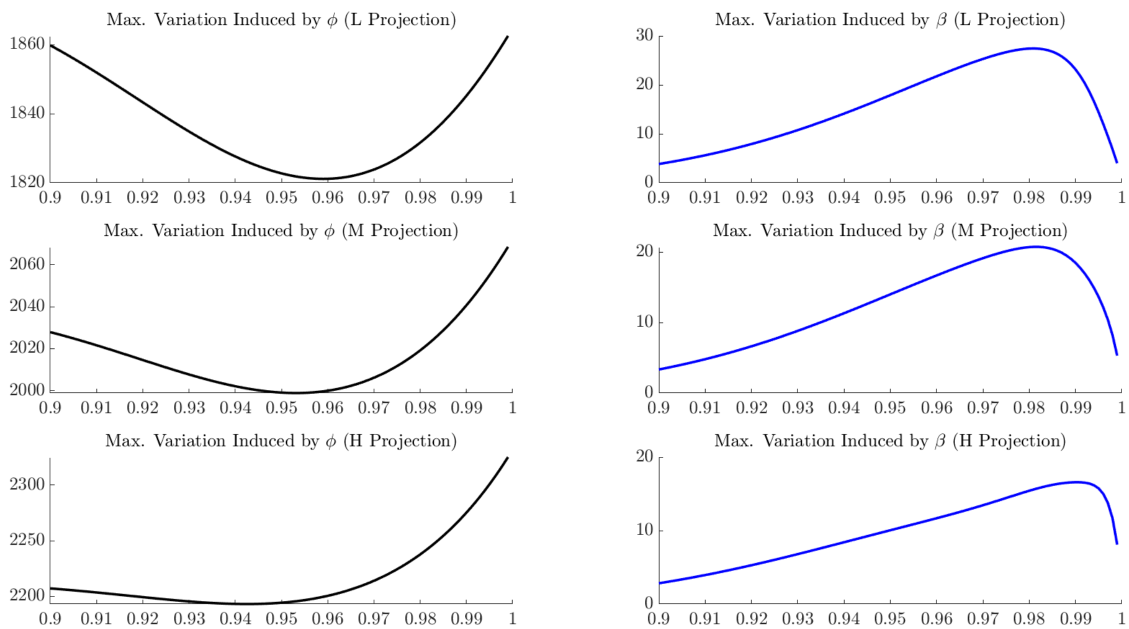

Notes: These figures illustrate how the optimal carbon tax changes as a result of variations in the inter- and intra-generational pure discount factors (ϕ and β, respectively). Graphs on the left-hand side show the percentage deviation in the optimal carbon tax between $\phi = 0.9$ and $\phi = 0.999$ as a function of β on the horizontal axis; graphs on the right-hand side show the percentage deviation in the optimal carbon tax between $\beta = 0.9$ and $\beta = 0.999$ as a function of ϕ on the horizontal axis. The three rows present the results using the low, medium and high variant United Nations (2019) global population projections. All values for ϕ and β are measured in annual equivalents; that is, the values used for the model results are given by the reported values raised to the power of 30.

As with Quaas and Bröcker (2016), I find that the optimal carbon tax is much more sensitive to assumptions on the intergenerational rate than the intragenerational rate: the maximum variation in the optimal carbon tax induced by changes in β is 30% compared to over 2000% for changes in ϕ. Changes in β and ϕ affect the optimal carbon tax in a number of ways. Both parameters affect the $\frac{\tilde{N}_s}{\tilde{N}_t}$ ratio in equation (1.36) by affecting how the different generations alive at a given point in time are weighted relative to each other in the policymaker's social welfare function. When β increases, this shifts the weight towards the old generation in $\tilde{N}_s$; when ϕ increases, this shifts the weight towards the young generation in $\tilde{N}_s$. Although the previous section of the chapter illustrated that the optimal carbon tax is sensitive to near-horizon changes in $\frac{\tilde{N}_s}{\tilde{N}_t}$, $\frac{\tilde{N}_s}{\tilde{N}_t}$ itself tends not to be sensitive to changes in β or ϕ given the forecasted patterns of demographic change. In particular, the asymmetries between the growth rate of the old generation over time relative to the young generation over time are not sufficiently large for this weighting effect to be significant in the context of this analysis. In addition to this channel, changes in β affect the optimal carbon tax via changes in savings rates, since higher values of β imply higher savings rates. What matters for the optimal carbon tax, however, is not the level of the savings rate but how it evolves over time. Conditional on the path of ψ_{s+1}^{s-1}, changes in β do not affect $\frac{(1-s_t)}{(1-s_s)}$ much. In contrast, ϕ directly affects the policymaker's intertemporal MRS so fluctuations in ϕ have large impacts on the optimal carbon tax.

This sensitivity analysis illustrates that normative judgements about how to weight the welfare of different generations are a key input into optimal carbon tax formulations. Optimal carbon tax results are likely to be highly sensitive to these normative judgements, and positive analysis is not necessarily a good substitute for such judgements.

1.5.2 The Fall in $r*$

In addition to the normative-positive division, the debate about discounting has recently been revisited in light of the decline in real interest rates that occurred following the Great Recession. Real interest rates on safe assets are now significantly lower than they were when the Nordhaus-Stern debate about discounting first emerged, with a number of authors arguing that optimal carbon tax estimates should therefore be revised upwards quite substantially (Bauer and Rudebusch, 2021; Carleton and Greenstone, 2021). My model shows that while these arguments are intuitively appealing, they are potentially misleading. Since both the real interest rate and the optimal carbon tax are endogenous objects that are simultaneously determined by the primitives of the model, it is not sufficient to simply revise the discount rates used in the optimal carbon tax formula while holding everything else constant. The underlying causes of the real interest rate decline need to be modelled and the equilibrium implications for the optimal carbon tax determined in a consistent manner.

My model accommodates two key mechanisms for the real interest rate decline: increased life expectancy and lower fertility rates. An increase in the life expectancy of retirees pushes up the savings rate, increasing the capital to labour ratio and pushing down returns. The decline in fertility rates reduces future labour supply further increasing the capital to labour ratio. While increases in the proportion of children surviving to middle age partially offsets the impact of lower fertility rates on labour supply, it does not completely offset it given the projected magnitudes of these changes. Even with an overall reduction in the future labour supply due to demographic change, there is an additional offsetting effect on the real rate that operates through the climate externality as a smaller labour force implies lower emissions. Less climate change in turn implies higher levels of productivity and thus higher returns to capital. Appendix A.4 formalises these results. For reasonable parameter values, however, the effect of higher life expectancies and lower birth rates pushes down real rates, as is the case in all calibration scenarios that I run.

While it is true that this real interest rate decline occurs while the optimal carbon tax is also rising in my model, this positive correlation is driven by the fact that my optimal carbon tax is a non-stationary object, since it is proportional to GDP. When analysing the optimal tax-to-GDP ratio, which is a stationary object, my model shows that the demographic transition simultaneously lowers the real interest rate and the optimal tax-to-GDP ratio. This implies that a decline in real interest rates caused by the demographic transition captured in my model should, in some sense, be associated with less aggressive climate policy than would have been the case absent the real rate decline. Figure 1.10 illustrates this point in two ways. First, for a given projection scenario, the optimal tax-to-GDP ratio is positively correlated with the marginal product of capital in the time series. Second, the optimal tax-to-GDP ratio is also positively correlated with the marginal product of capital across the various projection scenarios, with the scenarios that project a more drastic demographic transition delivering both lower returns to capital and a lower optimal tax-to-GDP ratio.

Figure 1.10: Marginal Product of Capital and Optimal Carbon Tax

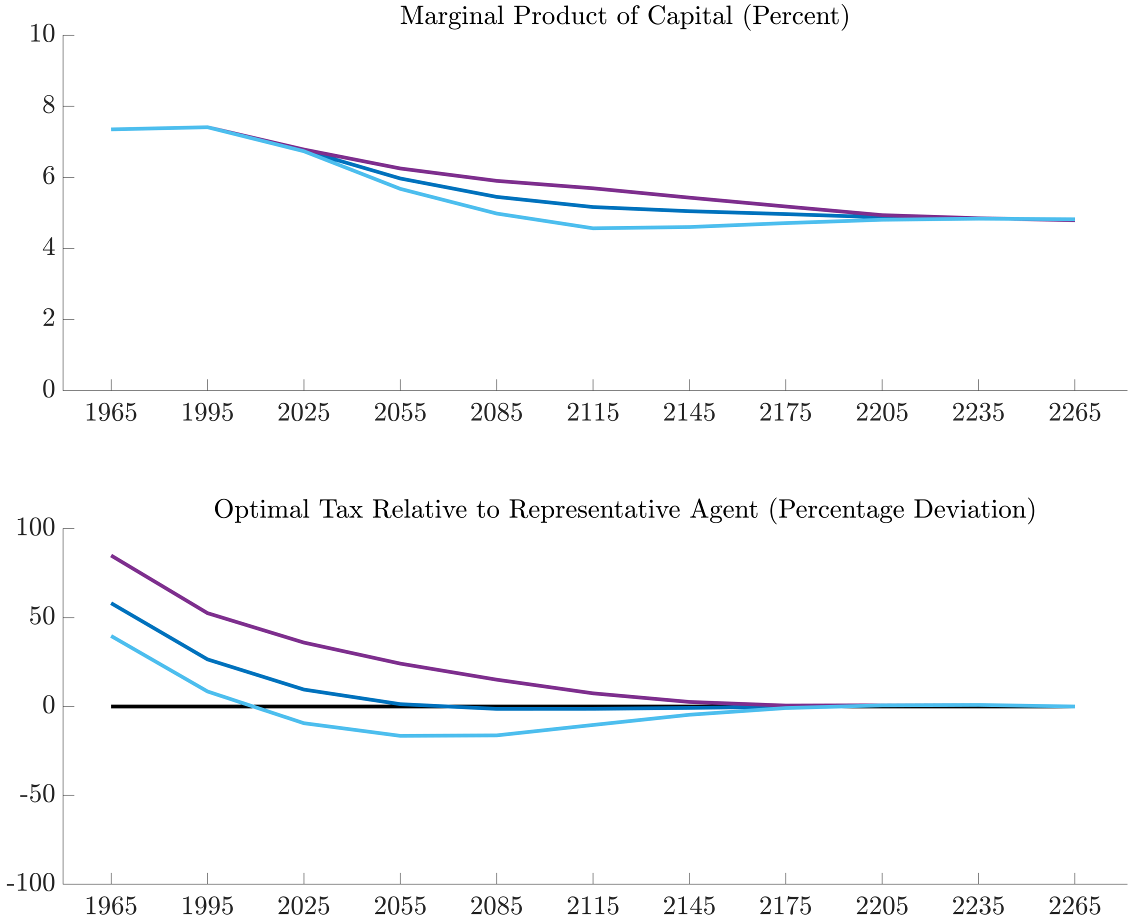

Notes: These figures illustrate how the marginal product of capital and the deviation between the optimal carbon tax and the representative agent equivalent evolve over time. The light blue, blue and purple lines present the results using the low, medium and high variant United Nations (2019) global population projections.

1.6 Conclusion

In this chapter, I introduce demographic change into an integrated assessment model of climate and the economy using a simple OLG structure. In doing so, I am able to derive in closed-form an expression for the optimal carbon-tax-to-GDP ratio. This ratio depends on a small number of parameters and exogenous variables, which can easily be pinned down from the data and existing literature. My formula shows in a transparent way how accounting for demographic change (in particular, changes in life expectancy and fertility rates) affects the optimal carbon tax. I then quantify these effects on the optimal carbon tax, and what the welfare effects of such policy adjustments are. Finally, I use my model to address some key debates in the literature regarding discounting and optimal carbon taxation.

The optimal carbon-tax-to-GDP ratio in my model depends on demographics via three channels. First, the policymaker's intertemporal marginal rate of substitution depends on the (adjusted) population growth rate, due to the concavity of the individual utility function. Second, this marginal rate of substitution also depends on changes in the aggregate savings rate, which are a function of life expectancy after retirement. Third, the optimal ratio factors in a second-best adjustment, which accounts for potentially suboptimal rates of capital accumulation introduced by the OLG structure of the model. The representative agent version of the model misses all three of these channels. By assuming log utility and an exponential damage function, the optimal tax ratio is not a function of other endogenous variables, such at the future level of output, which shuts down other demographic channels that could potentially influence the optimal carbon tax.

Quantitatively, accounting for the United Nations (2019) demographic projections changes the optimal carbon tax in the 2010-2040 period of the model by -9.3%, 9.5% and 36% relative to the representative agent benchmark under the low, medium and high variant population projections, respectively. These changes are largely driven by the effect of the adjusted population growth rate on the policymaker's intertemporal rate of substitution. The fact that this channel dominates the other two channels discussed in the previous paragraph is the consequence of two factors. First, the optimal tax ratio is proportional to this growth rate. Second, this growth rate is furthest from zero in the present before converging to zero over time. This pattern of convergence means that the largest effects in terms of deviations from the representative agent framework are the least heavily discounted. In contrast, the other two channels via which demographic change affects the optimal tax ratio are either mediated by non-linear functions that dampen their impact or are largest in the future when they are most heavily discounted. Consequently, the effect of the population growth rate dominates.

The welfare effects of accounting for these adjustments to the optimal carbon tax formula are also large. The gain in social welfare from adjusting the carbon tax such that is accounts optimally for demographic change is $626 Billion, $148 Billion, and $2.2 Trillion ($2015 lump-sum aggregate consumption equivalent) under the low, medium and high variant population projections, respectively. The welfare gains are the largest for the high variant projection for two reasons. First, the mistakes in setting the policy instrument by not accounting for demographic change are the largest. Second, a given deviation in the policy instrument from optimality is also more costly since climate change is a more acute issue with higher rates of population growth and thus higher emissions. While these welfare gains are large, it is worth noting that they are still small compared to the overall benefit of using carbon taxes (at most, the demographic adjustments account for 1.17% of the overall welfare gains associated with optimal carbon taxation).

Finally, I use my model to address debates regarding discounting and optimal carbon taxation. First, I show that when it comes to calibrating pure rates of time preference, it is important to differentiate between the policymaker's preferences and those of individual agents. In my model, the optimal carbon tax is very sensitive to the policymaker's pure rate of time preference but doesn't move much with individuals' pure rate of time preference. This suggests that optimal tax calculations will inherently have a strong normative component, even when we attempt to discipline our models with data. These issues are not always dealt with in a transparent manner in a representative agent setting, as it is natural in such models to assume that there is a single pure rate of time preference for agents and policymakers. Second, I show that reductions in the real interest rate do not automatically imply that climate policy should be more stringent. In fact, through the lens of my model, the observed demographic shifts over recent decades caused both a secular decline in real interest rates and also a lower optimal carbon-tax-to-GDP ratio.

Taken together, these findings highlight the importance of properly accounting for demographic change in IAMs that produce quantitative assessments of optimal climate policy. I have, however, made some strong assumptions on functional forms to incorporate key demographic mechanisms in a transparent and easily quantifiable manner. Future research is need to generalise these assumptions and add demographic change to richer IAMs. In doing so, the robustness of my quantitative results can be tested and new qualitative channels may also emerge. Although I do not explicitly address issues around uncertainty, the sensitivity of the results to the various population projections also highlights the fact that demographic uncertainty is a key source of uncertainty for climate policy that must be accounted for. Analysing this uncertainty in a more systematic way is another key area of future research. More generally, given the significant effects that the demographic transition and climate change will have on the global economy and the extent to which these two issues are intertwined, further research bringing together these two topics is much needed.

1.7 References

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of political economy*, 124(1):52–104.

Acemoglu, D. and Restrepo, P. (2022). Demographics and automation. *The Review of Economic Studies*, 89(1):1–44.

Aksoy, Y., Basso, H. S., Smith, R. P., and Grasl, T. (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, 11(1):193–222.

Attanasio, O., Kitao, S., and Violante, G. L. (2007). Global demographic trends and social security reform. *Journal of monetary Economics*, 54(1):144–198.

Barrage, L. (2018). Be careful what you calibrate for: social discounting in general equilibrium. *Journal of Public Economics*, 160:33–49.

Bauer, M. D. and Rudebusch, G. D. (2021). The rising cost of climate change: evidence from the bond market. *The Review of Economics and Statistics*, pages 1–45.

Belfiori, M. E. (2017). Carbon pricing, carbon sequestration and social discounting. *European Economic Review*, 96:1–17.

Belfiori, M. E. (2018). Climate change and intergenerational equity: Revisiting the uniform taxation principle on carbon energy inputs. *Energy Policy*, 121:292–299.

Bretschger, L. and Karydas, C. (2018). Optimum growth and carbon policies with lags in the climate system. *Environmental and Resource Economics*, 70(4):781–806.

Brock, W. A. and Hansen, L. P. (2018). Wrestling with uncertainty in climate economic models. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2019-71).

Cai, Y. and Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734.

Carleton, T. and Greenstone, M. (2021). Updating the united states government's social cost of carbon. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-04).

Carvalho, C., Ferrero, A., and Nechio, F. (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review*, 88:208–226.

Cipriani, G. P. (2014). Population aging and payg pensions in the olg model. *Journal of population economics*, 27(1):251–256.

Covert, T., Greenstone, M., and Knittel, C. R. (2016). Will we ever stop using fossil fuels? *Journal of Economic Perspectives*, 30(1):117–38.

Dietz, S. and Venmans, F. (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management*, 96:108–129.

Gerlagh, R., Jaimes, R., and Motavasseli, A. (2022). Global demographic change and climate policies. *Working Paper*.

Gerlagh, R. and Liski, M. (2018). Consistent climate policies. *Journal of the European Economic Association*, 16(1):1–44.

Gerlagh, R. and van der Zwaan, B. C. (2001). The effects of ageing and an environmental trust fund in an overlapping generations model on carbon emission reductions. *Ecological Economics*, 36(2):311–326.

Gertler, M. (1999). Government debt and social security in a life-cycle economy. In *Carnegie-Rochester Conference Series on Public Policy*, volume 50, pages 61–110. Elsevier.

Global Carbon Project (2021). Supplemental data of Global Carbon Project 2021 (1.0). *https://doi.org/10.18160/gcp-2021 Accessed on 31 July, 2022.*

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88.

Goulder, L. H. and Williams, R. C. (2012). The choice of discount rate for climate change policy evaluation. *Climate Change Economics*, 3(04):1250024.

Hassler, J., Krusell, P., and Smith Jr, A. A. (2016). Environmental macroeconomics. In *Handbook of macroeconomics*, volume 2, pages 1893–2008. Elsevier.

Howarth, R. B. (1996). Climate change and overlapping generations. *Contemporary Economic Policy*, 14(4):100–111.

Howarth, R. B. (1998). An overlapping generations model of climate-economy interactions. *Scandinavian Journal of Economics*, 100(3):575–591.

Howarth, R. B. (2000). Climate change and the representative agent. *Environmental and Resource Economics*, 15(2):135–148.

Kaplow, L., Moyer, E., and Weisbach, D. A. (2010). The social evaluation of intergenerational policies and its application to integrated assessment models of climate change. *The BE Journal of Economic Analysis & Policy*, 10(2).

Karp, L. and Rezai, A. (2014). The political economy of environmental policy with overlapping generations. *International Economic Review*, 55(3):711–733.

Kotlikoff, L., Kubler, F., Polbin, A., Sachs, J., and Scheidegger, S. (2021). Making carbon taxation a generational win win. *International Economic Review*, 62(1):3–46.

Krueger, D. and Ludwig, A. (2007). On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare. *Journal of monetary Economics*, 54(1):49–87.

Lemoine, D. and Rudik, I. (2017). Steering the climate system: using inertia to lower the cost of policy. *American Economic Review*, 107(10):2947–57.

Lemoine, D. and Traeger, C. (2014). Watch your step: optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–66.

Liang, J., Wang, H., and Lazear, E. P. (2018). Demographics and entrepreneurship. *Journal of Political Economy*, 126(S1):S140–S196.

Lisack, N., Sajedi, R., and Thwaites, G. (2021). Population aging and the macroeconomy. *International Journal of Central Banking*, 68.

Mier, M., Adelowo, J., and Weissbart, C. (2021). Taxation of carbon emissions and air pollution in intertemporal optimization frameworks with social and private discount rates. Technical report, ifo Working Paper.

Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the "dice" model. *The American Economic Review*, 83(2):313–317.

Nordhaus, W. D. (1994). *Managing the global commons: the economics of climate change*, volume 31. MIT press Cambridge, MA.

Nordhaus, W. D. (2007). To tax or not to tax: Alternative approaches to slowing global warming. *Review of Environmental Economics and policy*, 1(1).

Nordhaus, W. D. (2008). *A question of balance*. Yale University Press.

Nordhaus, W. D. and Boyer, J. (2003). *Warming the world: economic models of global warming*. MIT press.

Nordhaus, W. D. and Yang, Z. (1996). A regional dynamic general-equilibrium model of alternative climate-change strategies. *The American Economic Review*, 86(4):741–765.

Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, 99(2):281–290.

Piggott, J. and Woodland, A. (2016). *Handbook of the economics of population aging*. Elsevier.

Poterba, J. M. (2014). Retirement security in an aging population. *American Economic Review*, 104(5):1–30.

Quaas, M. F. and Bröcker, J. (2016). Substitutability and the social cost of carbon in a solvable growth model with irreversible climate change. Technical report, Economics Working Paper.

Rachel, Ł. and Summers, L. H. (2019). On secular stagnation in the industrialized world. *Brookings Papers on Economic Activity*, page 1.

Rezai, A. and van der Ploeg, F. (2016). Intergenerational inequality aversion, growth, and the role of damages: Occams rule for the global carbon tax. *Journal of the Association of Environmental and Resource Economists*, 3(2):493–522.

Rogner, H.-H. (1997). An assessment of world hydrocarbon resources. *Annual review of energy and the environment*, 22(1):217–262.

Stern, N. (2007). *The economics of climate change: the Stern review*. cambridge University press.

Traeger, C. P. (2015). Closed-form integrated assessment and uncertainty. *CESifo Working Paper Series*.

United Nations (2019). World population prospects 2019: Highlights. Technical report, United Nations, Department of Economic and Social Affairs, Population Division.

United States Environmental Protection Agency (2022). Greenhouse Gases Equivalencies Calculator - Calculations and References. *https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references Accessed on 22 April, 2022*.

Van den Bijgaart, I., Gerlagh, R., and Liski, M. (2016). A simple formula for the social cost of carbon. *Journal of Environmental Economics and Management*, 77:75–94.

von Below, D. (2012). Optimal carbon taxes with social and private discounting. In *SURED conference in Ascona, Switzerland*.

von Below, D., Dennig, F., and Jaakkola, N. (2013). Consuming more and polluting less today: Intergenerationally efficient climate policy. In *EAERE conference in Toulouse*, pages 26–29.

Wall Street Journal (2019). Economists' Statement on Carbon Dividends. *https://clcouncil.org/economists-statement/ Accessed on 3 February, 2022*.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The review of economics and statistics*, 91(1):1–19.

Williams, R. C., Gordon, H., Burtraw, D., Carbone, J. C., and Morgenstern, R. D. (2015). The initial incidence of a carbon tax across income groups. *National Tax Journal*, 68(1):195–213.

World Bank (2022). Carbon Pricing Dashboard. *https://carbonpricingdashboard.worldbank.org/ Accessed on 3 March, 2022.*

Chapter 2

Carbon Taxation in Small Open Economies

2.1 Introduction

Climate change is now widely recognised as a major policy issue, with the Paris Agreement being ratified by every country in the world except Eritrea, Iran, Libya and Yemen (United Nations Treaty Collection, 2022). Despite this global acceptance of the climate problem, the response remains fragmented. For example, of the 65 carbon pricing schemes that were in place as of April 2022, the vast majority operated at the national or sub-national level (World Bank, 2022). Aside from the EU Emissions Trading System (EU ETS), there is little international harmonisation of carbon pricing initiatives, and within most countries, carbon prices are not fully harmonised across sectors and locations. Despite this patchwork of carbon pricing initiatives, many models in the climate economics literature assume that climate policy takes the form of a globally harmonised carbon price. Making such an assumption fundamentally alters how optimal carbon pricing policies are formulated and what the welfare consequences of climate policy are. To better reflect the current policymaking environment, this chapter analyses the formulation and consequences of carbon pricing when such policies are implemented at the country-level.

In shifting to a country-level analysis, I make three key modelling decisions that reflect the policymaking reality faced by the typical economy: first, I assume a small, open economy; second, I impose a carbon budget on the economy; and third, I model a sectorally disaggregated economy. The first key assumption, that countries are small and open, means that all countries in my model take the path of climate change and international interest rates as given. Both of these assumptions are reasonable approximations for the vast majority of

Figure 2.1: Global GHG Emission Share by Country (2010-2018 Average)

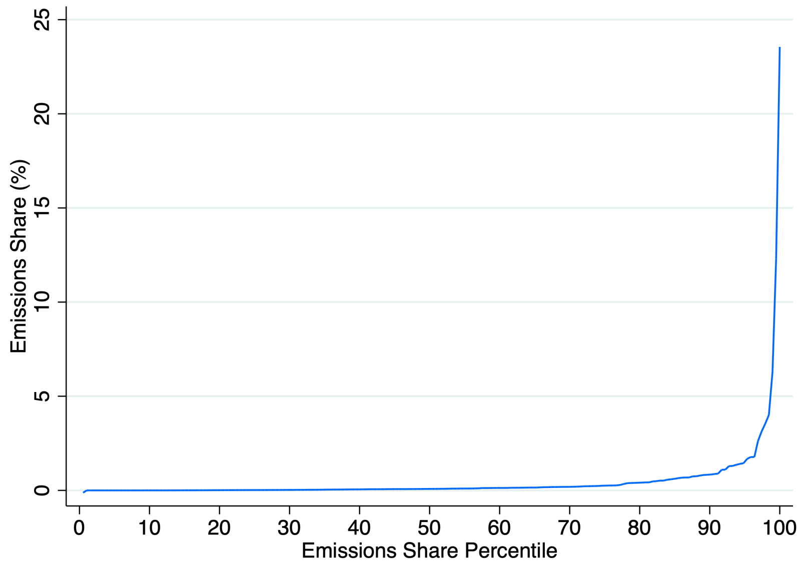

Notes: The the vertical axis of this figure measures the average share of global GHG emissions by country from 2010 to 2018, with countries organised by percentile on the horizontal axis. A total of 193 countries are included in the sample, accounting for just under 97% of the average global GHG emissions share over this period.
Source: Our World in Data (2022); Author's calculations.

Figure 2.2: Self-Imposed Climate Damages by Country (2010-2018 Average)

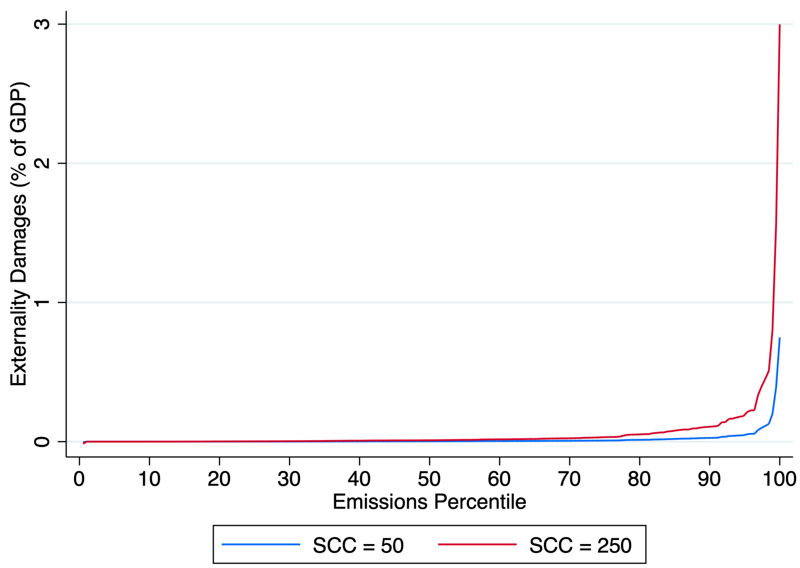

Notes: This figure illustrates the distribution of each country's self-imposed externality damages from their GHG emissions as a percentage of their GDP. These values reflect an average over the 2010-2018 period. The blue line shows the distribution assuming a SCC of $50; the red line shows the distribution assuming a SCC of $250 (measured in 2015 US Dollar terms).
Source: Our World in Data (2022); Author's calculations.

countries. Figure 2.1 shows the distribution of greenhouse gas (GHG) emissions across countries, which is highly skewed. Taking each country's average contribution to global emissions between 2010 and 2018, only countries above the 91st percentile of the distribution contributed more than 1 percent of global emissions. Aside from India, the United States and China, no country had an average contribution to global emissions above 5 percent during this period.

To estimate the cost of these country-level emissions in terms of self-inflicted output losses, I conduct the following back-of-the-envelope exercise. I multiply the average emissions for each country between 2010 and 2018 by two estimates of the social cost of carbon (SCC)[1] and divide this by the average value of global GDP over the same period. Assuming that climate-related GDP losses are evenly distributed and countries use a common discount rate, this figure provides an estimate of the discounted value of the current and future GDP losses a country imposes on itself via its emissions expressed as a percentage of its GDP. As shown by Figure 2.2, even assuming a SCC of \$250 per ton of CO_2, which is significantly higher than most SCC estimates, self-imposed climate damages are trivial for almost all countries: at the 90th percentile of this damage distribution, self-inflicted damages are still only 0.11 percent of GDP.

In addition to being small from a climate perspective, most countries are similarly small from an economic perspective, producing just a tiny fraction of global GDP. Given the correlation between emissions and output, the countries that tend to be the smallest emitters also tend to produce the least output (see Figure 2.3). Consequently, the joint assumption that countries take the path of climate change and global interest rates as given is a reasonable one in most cases.

[1] The SCC is the discounted value of all current and future global output losses associated with emitting a unit of carbon dioxide (CO_2) into the atmosphere.

Figure 2.3: Average Share of Global GDP and Emissions by Country (2010-2018)

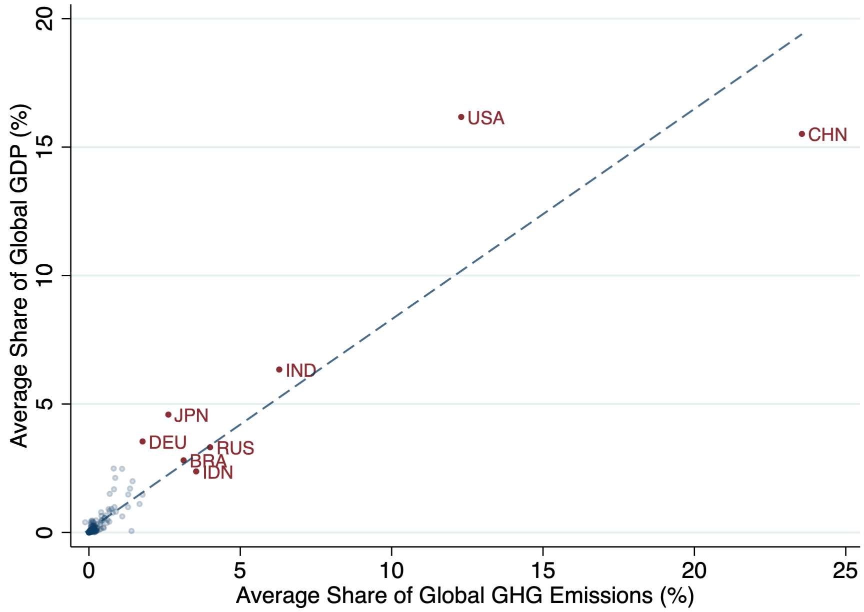

Notes: For a list of country codes used in this chapter, see Appendix B.1. Observations highlighted in red are those countries with GDP and Emission share contributions that are both above 1.5%.
Source: Our World in Data (2022); Author's calculations.

The second key modelling assumption I make is that countries are given an exogenously imposed emissions budget. In particular, I assume that each country has a fixed emissions budget from 2020 to 2049 that they cannot exceed, a period that I refer to in the rest of this chapter as the 'transition period'. The climate policy problem at the country-level is then how to allocate this budget over time. From 2050 onwards, I assume all countries become carbon neutral, again reflecting the typical policy preference observed in most countries' climate plans. I choose to impose an emissions budget for two reasons. From a modelling perspective, the existence of a climate-related constraint (either an emissions budget or a temperature limit) is necessary for a meaningful climate policy trade-off to emerge. In the global analysis of optimal carbon pricing, it is possible to formulate an unconstrained policy problem in which the policymaker must trade-off the desire to produce output (and thus emissions) in the present, with the need to preserve the climate to avoid output or utility losses in the future. For a country that takes the path of climate change as given, no such trade-off exists, and the optimal climate policy would simply be to do nothing absent any other market failures. Adding a binding emissions budget creates a different trade-off for the policy maker in which they need to manage the depletion of a finite resource that the market alone is unable to assign property rights to. In addition to being necessary from a modelling perspective, the existence of an emissions budget also reflects the reality of climate agreements. The main goal of the Paris Agreement is to limit global warming to well below 2 °C above pre-industrial levels (and ideally less that 1.5 °C). As shown in Table 2.1, the temperature limits specified by the Paris Agreement imply limits on how much more carbon the world can emit.

Table 2.1: Estimate of Global Carbon Budget (2020 Onwards, $GtCO_2$)

Temperature Limit	Likelihood of limiting global warming to temperature limit				
	17%	33%	50%	67%	83%
1.5 °C	900	650	500	400	300
1.7 °C	1450	1050	850	700	550
2.0 °C	2300	1700	1350	1150	900

Notes: The temperature limits are defined in terms of the maximum global warming allowed relative to the 1850–1900 global surface temperature. The budgets are expressed in terms of how much CO_2 can be emitted accounting for the expected global warming effect of non-CO_2 emissions.
Source: IPCC (2021).

It follows that for individual countries' policies to be consistent with the overall aim of the Paris Agreement, country-level carbon budgets also need to be met. While a number of countries have already put in place legally-binding carbon budgets[2], international disagreement on how the burden of climate policy should be shared remains a sticking point in the development of national-level policies consistent with the goals of the Paris Agreement. I side-step these issues by assuming that the global carbon budget is shared across countries on a per capita basis using population levels from 2020. There are, of course, many different methods that can be used to calculate what could be considered a fair distribution of carbon budgets across countries, and these can be easily incorporated into the calibration of the model presented in this chapter. There are also an interesting set of political economy questions about why the typical small country does not simply choose to free ride on the efforts of others when it comes to addressing climate change and how international climate agreements are formed, but that is not the focus of this chapter. Rather, I simply take it as given that in practice almost all countries in the world have now made pledges to reduce their GHG emissions as part of a collective effort to address climate change and I analyse the implications of these global-level agreements for country-level policymaking.

The third key modelling assumption that I make is to assume a sectorally disaggregated model of the economy. The main purpose of this disaggregation is to allow the model to

[2]The UK, for example, was the first country to set a legally binding carbon budget that limits the total amount of carbon that the UK can emit (UK Government, 2022). The budgets are set over a five year cycle with the first budget covering the 2008-2012 period. The UK's carbon budgets are agreed upon in advance and budgets are currently in place up to 2037.

accommodate sectorally targeted carbon policies. Most carbon pricing initiatives implemented to date impose different carbon prices across sectors or type of consumer. By using a sectorally disaggregated model, I am able to match existing policy regimes when calibrating the model using historical data and I can also assess the consequences of these sectorally targeted policies through my model simulations. Although I make the simplifying assumption in my quantitative exercises that all sectors are homogenous in the relative intensity with which they use dirty and green energy absent regulation, the model also accommodates sectoral heterogeneity in this dimension if needed.

In making these assumptions, I develop a tractable model of climate policy that can easily be calibrated and simulated using readily available data. The simplicity of the model also makes the relevant mechanisms transparent and easy to understand. The optimal policy requires that carbon taxes are equalised across all sectors of the economy at a given time and that these taxes grow over time at a rate equal to the real interest rate (the well-known Hotelling rule). The initial level of the optimal carbon tax can be expressed as the function of two reduced form objects: the ratio of laissez-faire emissions over the transition period relative to the carbon budget (hereafter referred to in short as the 'emissions ratio'), and the path of the laissez-faire dirty energy price. A higher emissions ratio implies a greater need for regulation to switch demand from dirty to green energy, requiring a greater shift in the relative price of these different energy varieties. Because energy varieties in my model are combined using a CES aggregator, this switching effect is a function of the percentage increase in the dirty energy price induced by the introduction of the carbon tax. Consequently, the higher the laissez-faire cost of producing dirty energy, the larger the carbon tax needs to be to achieve a given level of demand switching away from dirty energy.

Under my baseline calibration assumptions (notably, a real rate of 2%, a carbon budget consistent with limiting global temperature change to less than 1.7 °C, and an elasticity of substitution between dirty and green energy of 1.5), I find an average optimal carbon tax value of $325 per ton of CO_2 in 2020 for the 35 countries in my sample[3]. The average value of the discounted output losses over the transition period associated with the optimal policy is 2.1% relative to the no policy scenario. There is considerable variation in these results across countries, however, with the optimal 2020 carbon tax values ranging from a low of $4 per ton of CO_2 in Malta to a high of $749 per ton of CO_2 in Taiwan. As the carbon budget is relaxed and the elasticity of substitution across green and dirty energy is increased, the average optimal carbon tax drops significantly at all points in time. In contrast, although the value of the optimal carbon tax drops in the early part of the transition period as the

[3]This value is expressed in 2017 USD PPP terms, as are all other monetary values in this chapter unless otherwise specified.

real interest rate is raised, there is a large increase in future carbon taxes at some point along the transition path. For the sample of countries studied, the optimal path of carbon taxation is not sensitive to expected uncertainty around future labour force projections. This insensitivity follows from the fact that for most of the countries in my sample, there is little variation across the main United Nations (2019) demographic projections, with the average working age population across countries expected to shrink slightly from 2020-2049.

The average output loss associated with the optimal policy moves in the same direction as the optimal carbon tax given changes in the carbon budget, elasticity of substitution and population projections. In contrast, the average output loss is not sensitive to the real interest rate. This follows from the fact that the carbon tax in any given period reflects the marginal output loss from climate policy in that period. Changes in the carbon budget, elasticity of substitution and labour force projections change the cost of regulation, and thus shift the entire path of carbon tax in the same direction as this change in costs is optimally smoothed over the transition period. Changes in the real rate, on the other hand, do not affect the costs of regulation at any given point, but rather affect how the policymaker wishes to smooth these costs intertemporally. Thus, while the optimal tax value in 2020 is highly sensitive to the path of real rates over the transition period, the welfare costs of policy are not.

In addition to studying the optimal policy across countries, I examine suboptimal policies involving sectoral exemptions. Most carbon pricing policies provide exemptions to certain sectors of the economy. I show that under a set of simplifying assumptions (in particular, that each sector's demand for dirty energy relative to green energy is symmetric absent any taxes), the welfare costs of providing sectoral tax exemptions are a function of the proportion of energy expenditures accounted for by the exempt sectors. I also show that providing full tax exemptions to a subset of sectors places an upper bound on the emissions reductions that can be achieved relative to the laissez-faire equilibrium. This upper bound is equal to the proportion of energy expenditures accounted for by the non-exempt sectors.

Quantitatively, I find that under the carbon budgets consistent with the 1.5 °C, 1.7 °C and 2.0 °C temperature limits, the maximum share of energy expenditures that can be feasibly exempt is 25.1%, 46.0% and 70.0%, respectively, on average across my sample. To put these numbers into context, the average share of energy expenditure accounted for by transport, manufacturing and final demand by households in my sample is 12.5%, 22.8% and 36.3%, respectively, with the other nine non-energy sectors I aggregate the economy into accounting for the remaining 28.3% of energy expenditures in a manner that is fairly evenly distributed.

I then use my model to compute the tax implications and output losses when sectoral exemptions to the carbon tax are granted but the tax remains intertemporally optimised. Under the baseline calibration scenario, the average path of carbon taxes increases by 59.6% when 40% of the maximum possible exemptions are granted by a country and 333.1% when 80% of the maximum possible exemptions are granted by a country. In the limit, carbon taxes approach infinity as 100% of possible exemptions are granted. The tax increase necessitated by providing such sectoral exemptions is decreasing in the elasticity of substitution between dirty and green energy as the non-exempt sectors of the economy are more responsive the higher tax rates.

Turning to the welfare implications, 0.8% and 2.9% of discounted GDP over the transition period is lost relative to the optimal policy when 40% and 80% of the maximum exemptions are granted, respectively. The output losses associated with the sectorally inefficient policy are higher when either the carbon budget or elasticity of substitution are lower. This follows from the fact that a lower carbon budget or elasticity of substitution increases the average gap between the optimal tax and the suboptimal tax in each sector, thus increasing the distortions associated with the suboptimal tax policy.

These results are relevant to policy in that they provide guidance on how one of the main climate policy instruments ought to be set. They also speak to the politically sensitive issue of the welfare costs of such policies, and how the common practice of providing sectoral exemptions to carbon taxation affects these welfare costs. The results suggest that the output losses associated with meeting the targets of the Paris Agreement are relatively small under most assumptions. The inefficiency cost of providing sectoral exemptions is initially trivial, but as exemptions are increased towards their maximum level, the inefficiency costs eventually become larger than the cost of the optimal policy itself.

Related Literature. This chapter contributes the extensive economics literature on optimal carbon pricing (see for example, Barrage, 2020; Cai and Lontzek, 2019; Dietz et al., 2018; Dietz and Venmans, 2019; Golosov et al., 2014; Lemoine and Rudik, 2017; Nordhaus, 2008; Stern, 2007; van der Ploeg, 2018). This literature has tended to focus on optimal carbon pricing at the global-level absent any temperature or emissions constraints. Rather, the optimal path of emissions (and associated carbon price) is solved to maximise global welfare, irrespective of whether such a policy is consistent with the temperature constraints stipulated by global climate agreements. It is often the case that optimal policies determined within the context of these models overshoot the temperature limits imposed by international agreements (Tol, 2013), limiting the policy applications of these results in a world where countries have committed to temperature limits and emissions constraints.

While there are a number of studies that assess optimal carbon pricing under a temperature or carbon budget constraint (Dietz and Venmans, 2019; Emmerling et al., 2019;

Gollier, 2021; Lemoine and Rudik, 2017; Olijslagers et al., 2021; van der Ploeg, 2018), these studies fail to incorporate all of my other key assumptions: that is, they do not take the path of interest rates and climate change as given while also assuming a sectorally disaggregated economy. Sectorally disaggregated modelling, in particular, tends to be rare in optimal carbon pricing studies within the economics literature. While there are a number of papers that focus on carbon pricing in disaggregated economies, these tend not to study policy formulation. Instead, they examine the impact on the economy of imposing some predetermined policy (Cavalcanti et al., 2021; Devulder and Lisack, 2020; Frankovic, 2022; King et al., 2019). The climate policy studies that do adopt all of my main assumptions to derive optimal policy paths that are consistent with the temperature limits in global climate agreements tend not to use the same fully specified macroeconomic models with competitive markets found in the economics literature (Clarke et al., 2014). The aim of this chapter is, therefore, to bridge this gap by developing a tractable, policy-relevant macroeconomic model that can be easily taken to the data for a wide range of countries to provide guidance on how carbon pricing policies should be implemented across sectors and over time.

Roadmap. The rest of the chapter is organised as follows. Section 2 outlines the model. Section 3 characterises the competitive equilibrium. Section 4 of the chapter examines the optimal carbon tax and welfare implications of this policy across a number of countries. Section 5 uses the model to assess the tax and welfare implications of suboptimal policy in the form of sectoral tax exemptions when the intertemporal aspect of policy remains optimised. Section 6 assess the fit of the model with the data and what this implies for the results. Section 7 concludes.

2.2 Model Setup

This section of the chapter outlines my model for a given small, open economy. The economy is made up of N non-energy sectors and one energy sector that has dirty and green energy varieties. Firms in the non-energy sectors produce output using capital, labour, energy and intermediate inputs from other non-energy sectors. The economy is able to freely trade an aggregate final good that is the composite of the $N + 1$ sectoral goods, and faces an exogenously determined interest rate on foreign borrowing and lending. The government levies carbon taxes on dirty energy consumption and rebates the proceeds lump sum to households. For convenience, no notation is used to index the economy under consideration, although the quantitative analysis simulates the model outlined in this section across a number of different economies.

Although this paper focuses on sectoral heterogeneity in climate policy, I setup a version of the model here that also accommodates heterogeneity in each sector's ability to use and substitute between dirty and green energy. I shut down this heterogeneity in the model for the purposes of calibration, but I illustrate how it can be incorporated as this is a potentially important extension to my quantitative results.

2.2.1 Households

The economy is populated by an infinitely lived representative agent with lifetime utility in period t given by

$$W_t = \sum_{s=t}^{\infty} \beta^{s-t} U\left(C_s\right) \tag{2.1}$$

where C_t is the aggregate consumption good and $U(.)$ is an instantaneous utility function with the standard properties[4]. The agent is endowed with L_t units of labour in each period t, which are supplied inelastically. The period budget constraint of the agent, specified in real terms, is given by

$$C_t + K_{t+1} + B_{t+1} = w_t L_t + r_{k,t} K_t + \pi_t + T_t + (1-\delta) K_t + (1 + r_{b,t}) B_t \tag{2.2}$$

w_t is the agent's wage, K_t is the agent's capital holdings at time t, $r_{k,t}$ is the rental rate for capital paid by firms, and δ is the depreciation rate. B_t are bond holdings at time t, and $r_{b,t}$ is the exogenously determined return on these bonds. π_t are aggregate profits in the economy and T_t are lump sum transfers from the government to households.

2.2.2 Firms

Although the aggregate consumption good is aggregated via the representative agent's utility function, I assume that a firm exists that transforms output from various sectors into a scalar measure of aggregate output via a production function. This assumption has no bearing on the results but is notationally convenient. In particular, I assume that the aggregate output producer has a production function given by:

$$Y_t = \left(\frac{Y_{e,t}}{\omega_{ye}}\right)^{\omega_{ye}} \left[\prod_{i=1}^{N} \left(\frac{Y_{i,t}}{\omega_{yi}}\right)^{\omega_{yi}}\right] \tag{2.3}$$

where $Y_{i,t}$ represent final goods from sector i, and $Y_{e,t}$ are final goods from the energy sector. Abusing notation slightly, let N denote both the number of non-energy sectors and the set

[4] I assume $U'(.) > 0$ and $U''(.) < 0$.

of non-energy sectors: $i \in \{1, 2, ..., N\}$. The N sectoral producers have production functions given by

$$Q_{i,t} = A_{i,t} \left(\frac{L_{i,t}}{\omega_{il}}\right)^{\omega_{il}} \left(\frac{K_{i,t}}{\omega_{ik}}\right)^{\omega_{ik}} \left(\frac{Z_{ie,t}}{\omega_{ie}}\right)^{\omega_{ie}} \left[\prod_{j=1}^{N} \left(\frac{Z_{ij,t}}{\omega_{ij}}\right)^{\omega_{ij}}\right] \ \forall\ i \in N \tag{2.4}$$

where $L_{i,t}$ is their labour input, $K_{i,t}$ is their capital input, $Z_{ie,t}$ is sector i's use of energy as an intermediate input, $Z_{ij,t}$ is sector i's use of intermediate inputs from sector j and $A_{i,t}$ is the exogenously determined sectoral TFP. Defining the set N_y as $N_y \equiv N \cup \{y\}$, the production function for the sector $i \in N_y$ energy composite is given by

$$Q_{ie,t} = A_{ie,t} \left(\omega_{id} Q_{id,t}^{\frac{\sigma_{ie}-1}{\sigma_{ie}}} + \omega_{ig} Q_{ig,t}^{\frac{\sigma_{ie}-1}{\sigma_{ie}}}\right)^{\frac{\sigma_{ie}}{\sigma_{ie}-1}} \ \forall\ i \in N_y \tag{2.5}$$

where $Q_{id,t}$ and $Q_{ig,t}$ are sector i's respective use of dirty and green energy inputs, $A_{ie,t}$ is exogenously determined energy efficiency, and σ_{ie} is a sector-specific elasticity of substitution. The dirty and green energy varieties are produced linearly from labour:

$$Q_{ik,t} = A_{ik,t} L_{ik,t} \ k \in \{d, g\}, \ \forall\ i \in N_y \tag{2.6}$$

$A_{ik,t}$ are exogenously determined energy sector productivities, and the units of dirty energy are normalised such that one unit of energy produces one unit of CO_2 emissions and one unit of green energy provides an equivalent amount of physical energy output as one unit of dirty energy. The market clearing conditions for labour, capital, sectoral goods and the energy market, respectively, are given by

$$L_t = \sum_{i=1}^{N} L_{i,t} + \sum_{k \in \{d,g\}} \sum_{i \in N_y} L_{ik,t} \tag{2.7}$$

$$K_t = \sum_{i=1}^{N} K_{i,t} \tag{2.8}$$

$$Q_{i,t} = Y_{i,t} + \sum_{j=1}^{N} Z_{ji,t} \ \forall\ i \in N \tag{2.9}$$

$$Q_{ie,t} = Z_{ie,t} \ \forall\ i \in N \tag{2.10}$$

$$Q_{ye,t} = Y_{e,t} \tag{2.11}$$

The market clearing conditions combined with the capital accumulation equation $K_{t+1} = I_t + (1 - \delta) K_t$ also imply the aggregate goods market clearing condition

$$Y_t = C_t + I_t + NX_t \tag{2.12}$$

where NX_t are net exports of the final good, which automatically adjust to ensure that (2.12) holds. All markets are assumed to be perfectly competitive.

2.2.3 Government

The government is able to set carbon taxes on each sector $i \in N_y$ (denoted by $\tau_{i,t}$), which absent uncertainty within a given period, is equivalent to setting the quantity of emissions for each sector directly via a set of segmented emissions trading schemes. The government returns the proceeds from these taxes to households in the form of a lump-sum rebate T_t. The government's policies must be such that it runs a balanced budget in all periods:

$$\sum_{i \in N_y} \tau_{i,t} Q_{id,t} = T_t \tag{2.13}$$

In addition to the balanced budget constraint, the government is also required to ensure that from some period s to T, cumulative carbon dioxide emissions do not exceed a fixed amount $\bar{Q}_d$ (which denotes the carbon budget):

$$\sum_{t=s}^{T} \sum_{i \in N_y} Q_{id,t} \leq \bar{Q}_d \tag{2.14}$$

2.3 Competitive Equilibrium

This section of the chapter defines the competitive equilibrium of the economy and illustrates how the competitive equilibrium of the economy is solved conditional on the carbon tax policy in place. In addition, the response of emissions to sectorally-targeted tax changes is derived. In particular, total emissions in the economy are decreasing in the value of the carbon tax in any given sector of the economy, but there is a lower bound on the emissions that can be achieved using carbon taxes on just a single sector. This lower bound is given by the laissez-faire emissions of all sectors that face no carbon tax.

2.3.1 Definition

A competitive equilibrium in this economy is formally defined as follows

Definition 1: A competitive equilibrium consists of a sequence of allocations
$\Big\{C_t, K_{t+1}, I_t, Y_t, NX_t, B_{t+1}, \{Y_{i,t}, Q_{ie,t}, Q_{id,t}, Q_{ig,t}, L_{id,t}, L_{ig,t}\}_{i \in N_y},$
$\Big\{Q_{i,t}, L_{i,t}, K_{i,t}, Z_{ie,t}, \{Z_{ij,t}\}_{j \in N}\Big\}_{i \in N}\Big\},$

prices $\left\{r_{b,t}, r_{k,t}, w_t, \{p_{ie,t}, p_{id,t}, p_{ig,t}\}_{i\in N_y}, \{p_{i,t}\}_{i\in N}\right\}$, and policies $\left\{\{\tau_{i,t}\}_{i\in N_y}, T_t\right\}$ such that:

1. the allocations solve the consumers' utility maximisation problems and the firms' profit maximisation problems given prices and polices,
2. the government budget constraint is satisfied in every period,
3. capital accumulates according to the capital accumulation equation, and
4. markets clear.

A competitive equilibrium that is consistent with the country's climate targets is one in which conditions 1-4 hold and the carbon budget (2.14) is also satisfied.

2.3.2 Prices

The no arbitrage condition on capital and bonds implies $r_{k,t}$ is pinned down by the exogenously determined real rate on bonds and capital depreciation:

$$r_{k,t} = r_{b,t} + \delta \tag{2.15}$$

Given unit cost pricing and the tax policy, the price of dirty energy supplied to sector i is given by

$$p_{id,t} = \frac{w_t}{A_{id,t}} + \tau_{i,t} \ \forall\ i \in N_y \tag{2.16}$$

Defining $\phi_{i,t} \equiv 1 + \tau_{i,t}\left(\frac{w_t}{A_{d,t}}\right)^{-1}$, the dirty energy price for sector i can be expressed as

$$p_{id,t} = \frac{w_t}{A_{id,t}}\phi_{i,t} \ \forall\ i \in N_y \tag{2.17}$$

$\phi_{i,t}$ is therefore the sector i dirty energy price ratio with and without the tax. As shown later in this section, $\phi_{i,t}$ is also a sufficient statistic for measuring the stringency of tax policy. This follows from the fact that energy varieties are aggregated using a CES function. The green energy price in sector i is given by

$$p_{ig,t} = \frac{w_t}{A_{ig,t}} \ \forall\ i \in N_y \tag{2.18}$$

Standard CES properties imply that the composite energy price in sector i is given by

$$p_{ie,t} = A_{ie,t}^{-1}\left(\omega_{id}^{\sigma_{ie}} p_{id,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} p_{g,t}^{1-\sigma_{ie}}\right)^{\frac{1}{1-\sigma_{ie}}} \ \forall\ i \in N_y \tag{2.19}$$

Substituting (2.17) and (2.18) into (2.19) then implies

$$p_{ie,t} = A_{ie,t}^{-1} w_t \left(\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1} \right)^{\frac{1}{1-\sigma_{ie}}} \ \forall \ i \in N_y \tag{2.20}$$

Cobb-Douglas sectoral production gives the following sectoral output prices:

$$p_{i,t} = A_{i,t}^{-1} w_t^{\omega_{il}} r_{k,t}^{\omega_{ik}} p_{ie,t}^{\omega_{ie}} \left(\prod_{j=1}^{N} p_{j,t}^{\omega_{ij}} \right) \ \forall \ i \in N \tag{2.21}$$

I next define the standard Leontief Inverse $\psi \equiv [I_N - \Omega]^{-1}$, where I_N is the $N \times N$ identity matrix and Ω is the $N \times N$ matrix with $\Omega_{ij} = \omega_{ij}$. As shown in Appendix B.2, solving the system of equations in (2.21) implies

$$p_{i,t} = \left(\prod_{j=1}^{N} A_{j,t}^{-\psi_{ij}} \right) w_t^{\left(\sum_{j=1}^{N} \psi_{ij} \omega_{jl} \right)} r_{k,t}^{\left(\sum_{j=1}^{N} \psi_{ij} \omega_{jk} \right)} \left(\prod_{j=1}^{N} p_{je,t}^{\psi_{ij} \omega_{je}} \right) \ \forall \ i \in N \tag{2.22}$$

The aggregate price index (which will be normalised to one) is then given by

$$p_{y,t} = p_{ye,t}^{\omega_{ye}} \left(\prod_{i=1}^{N} p_{i,t}^{\omega_{yi}} \right) \tag{2.23}$$

Before deriving an expression for $p_{y,t}$ that implicitly pins down w_t, it is useful to introduce some additional notation. The sales share of each sector is defined as

$$\lambda_{i,t} \equiv \frac{p_{i,t} Q_{i,t}}{Y_t} \ \forall \ i \in N \tag{2.24}$$

$$\lambda_y \equiv 1 \tag{2.25}$$

Given the Cobb-Douglas functional forms and market clearing conditions, the (time invariant) equilibrium value of λ_i is given by (see Appendix B.3 for details)

$$\lambda_{i,t} = \sum_{j=1}^{N} \omega_{yj} \psi_{ji} \ \forall \ i \in N \tag{2.26}$$

I then define the non-energy labour share, capital share, and energy shares as

$$\lambda_{l,t} \equiv \frac{\sum\limits_{i=1}^{N} w_t L_{i,t}}{Y_t} \tag{2.27}$$

$$\lambda_{k,t} \equiv \frac{\sum\limits_{i=1}^{N} r_{k,t} K_{i,t}}{Y_t} \tag{2.28}$$

$$\lambda_{ie} \equiv \frac{p_{ie,t} Q_{ie,t}}{Y_t} \tag{2.29}$$

$$\lambda_e \equiv \sum_{i \in N_y} \lambda_{ie} \tag{2.30}$$

Again, these are time invariant objects in equilibrium given by (see Appendix B.3 for details)

$$\lambda_l = \sum_{i=1}^{N} \lambda_i \omega_{il} \tag{2.31}$$

$$\lambda_k = \sum_{i=1}^{N} \lambda_i \omega_{ik} \tag{2.32}$$

$$\lambda_{ie} = \lambda_i \omega_{ie} \tag{2.33}$$

By combining the household budget constraint and aggregate goods market clearing condition, it also follows that (see Appendix B.3 for details)

$$\lambda_l + \lambda_k + \lambda_e = 1 \tag{2.34}$$

Combining (2.20) and (2.22) with (2.23) then implies

$$p_{y,t} = A_t^{-1} w_t^{1-\lambda_k} r_{k,t}^{\lambda_k} \left[\prod_{i \in N_y} \left(\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1} \right)^{\frac{\lambda_{ie}}{1-\sigma_{ie}}} \right] \tag{2.35}$$

where $A_t \equiv \left(\prod\limits_{i=1}^{N} A_{i,t}^{\lambda_i} A_{ie,t}^{\lambda_{ie}} \right)$ is a measure of aggregate TFP (excluding TFP of the dirty and green energy varieties).

Proposition 1: In any period t and for any set of carbon taxes $\{\tau_{i,t}\}_{i \in N_y}$, there always exists a unique set of prices $\left\{ r_{b,t}, r_{k,t}, w_t, \{p_{ie,t}, p_{id,t}, p_{ig,t}\}_{i \in N_y}, \{p_{i,t}\}_{i \in N} \right\}$ that satisfy the competitive equilibrium conditions.

Proof: Conditional on w_t and $\{\tau_{i,t}\}_{i \in N_y}$, it has been shown that all prices exist and are uniquely determined. It thus suffices to show that a unique w_t exists that satisfies the competitive equilibrium conditions. Defining the function

$$F_{p,t}\left(w_t\right) \equiv A_t^{-1} w_t^{1-\lambda_k} r_{k,t}^{\lambda_k} \left[\prod_{i \in N_y} \left(\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1} \right)^{\frac{\lambda_{ie}}{1-\sigma_{ie}}} \right] \tag{2.36}$$

it follows that w_t is implicitly determined from the normalising condition $F_{p,t}\left(w_t\right) = 1$. Since

$$\frac{dF_{p,t}\left(w_t\right)}{dw_t} > 0$$

$$F_{p,t}\left(0\right) = 0$$

$$\lim_{w_t \to \infty} F_{p,t}\left(w_t\right) = \infty$$

there always exists a unique w_t^* such that $F_{p,t}\left(w_t^*\right) = 1$ holds.

2.3.3 Quantities

Conditional on prices, all quantities can be determined in closed form. The full determination of equilibrium quantities is outlined in Appendix B.4; I focus here on the determination of dirty energy inputs $Q_{id,t}$ since this is the key quantity that the policymaker must control when setting carbon taxes. Given the CES production function employed to produce the composite energy good, these energy producers have dirty and green energy expenditure shares given by

$$p_{id,t} Q_{id,t} = \alpha_{id,t} p_{ie,t} Q_{ie,t} \;\forall\; i \in N_y \tag{2.37}$$

$$p_{ig,t} Q_{ig,t} = \alpha_{ig,t} p_{ie,t} Q_{ie,t} \;\forall\; i \in N_y \tag{2.38}$$

where

$$\alpha_{id,t} \equiv \frac{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}}}{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}} \;\forall\; i \in N_y \tag{2.39}$$

$$\alpha_{ig,t} \equiv \frac{\omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}}{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}} \;\forall\; i \in N_y \tag{2.40}$$

All revenue generated by the green energy sector goes to labour, while the share of $p_{id,t}Q_{id,t}$ paid to $L_{id,t}$ is $\frac{1}{\phi_{i,t}} \leq 1$. Together, these facts imply that the labour allocations are given by

$$L_{id,t} = \left(\frac{\frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}}{\lambda_l + \sum\limits_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum\limits_{j \in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall \ i \in N_y \tag{2.41}$$

$$L_{ig,t} = \left(\frac{\lambda_{ie}\alpha_{ig,t}}{\lambda_l + \sum\limits_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum\limits_{j \in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall \ i \in N_y \tag{2.42}$$

$$L_{i,t} = \left(\frac{\lambda_{i}\omega_{il}}{\lambda_l + \sum\limits_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum\limits_{j \in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall \ i \in N_y \tag{2.43}$$

Dirty energy production is therefore given by

$$Q_{d,t} = \sum_{i \in N_y} Q_{id,t} = \left(\frac{\sum\limits_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}} A_{id,t}}{\lambda_l + \sum\limits_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}} + \sum\limits_{i \in N_y} \lambda_{ie}\alpha_{ig,t}} \right) L_t \tag{2.44}$$

Proposition 2: If $\sigma_{ie} > 1$ and $A_{id} = A_d \ \forall \ i \in N_y$,

$$\frac{dQ_{d,t}}{d\tau_{i,t}} < 0$$

$$\lim_{\tau_{i,t} \to \infty} Q_{d,t} = \left(\frac{\sum\limits_{j \in N_0} \lambda_{je}\alpha_{jd,t}A_{jd,t}}{\lambda_l + \lambda_e} \right) L_t$$

Proof: See Appendix B.5.

Proposition 2 states that under the assumption that all elasticities of substitution between dirty and green energy are greater than one and dirty energy TFP is the same across sectors, raising carbon taxes in any single sector reduces economy-wide emissions. For the results in the subsequent sections of the chapter, it is useful to understand the mechanisms underlying this statement. Raising $\tau_{i,t}$ leads to an increase in $\phi_{i,t}$ (the ratio of the dirty energy price with the tax relative to the price without the tax). This happens via the direct effect of increasing $\tau_{i,t}$ and also the indirect effect of lowering real wages w_t, both of which increase the size of the tax faced by sector i relative to the cost of producing dirty energy. This causes an increase in the price of dirty dirty energy relative to green energy for sector i, leading to a substitution away from dirty energy. Although sectors $j \neq i$ do not see a direct increase in the carbon tax they face, because their tax $\tau_{j,t}$ is held constant in real terms, the general equilibrium effect of lower real wages also raises the value of taxes in these sectors relative to the cost of producing dirty energy provided $\tau_{j,t} > 0$. This relative price increase also leads these sectors to substitute away from dirty energy usage despite not being directly affected by the carbon tax increase. Since these indirect substitution effects only operate when $\tau_{j,t} > 0$, it follows that the lower bound on emissions that can be achieved by raising carbon taxes on one sector alone and holding taxes constant in real terms on all other sectors is given by the total emissions in the laissez-faire equilibrium of all sectors facing a zero tax.

2.4 Optimal Policy

This section outlines the optimal carbon tax policy that satisfies the economy's carbon budget constraint. This policy is optimised across sectors (a static efficiency condition) and over time (an intertemporal efficiency condition). The static efficiency condition dictates that a common carbon tax should be applied to all sectors of the economy in a given period. The intuition for this results follows from the fact that the price of emissions in a given sector ($\tau_{i,t}$) is equal to the marginal change in aggregate output from increasing emissions in that sector. It follows that conditional on an aggregate emissions budget in a given period, output is maximised when $\tau_{i,t}$ is equalised across all sectors. Similarly, the carbon tax applied to a given period (τ_t) is equal to the marginal effect on output from increasing aggregate emissions in that period. Given that emissions are a scarce and finite resource under a binding carbon constraint, the Hotelling rule applies: the carbon tax should grow at a rate equal to the real interest rate on international borrowing to maximise the discounted value of the economy's aggregate output over the transition period. This section first outlines the policy problem and formally derives the efficiency conditions that characterise the optimal policy. I then discuss how the model is calibrated to produce quantitative estimates of the optimal carbon tax across a sample of 35 economies before turning to a presentation of these results.

2.4.1 Optimal Policy Problem

Given the small open economy assumption, maximising individual welfare over the transition period is equivalent to maximising the present discounted value of GDP over this period. The policymaker's problem can therefore be written as

$$\max_{\left\{\{\tau_{i,t}\}_{i\in N_y}\right\}_{t=1}^{T}} \quad \sum_{t=1}^{\infty} \frac{Y_t\left(\{\tau_{i,t}\}_{i\in N_y}\right)}{R_t}$$

$$\text{s.t.} \sum_{t=1}^{T} \sum_{i \in N_y} \bar{Q}_{id,t} \leq \bar{Q}_d$$

where I define

$$R_t \equiv \prod_{s=2}^{t} (1 + r_{b,s})$$

Proposition 3:

A unique optimal policy $\left\{\left\{\tau_{i,t}^*\right\}_{i \in N_y}\right\}_{t=1}^{T}$ always exists. When the emissions constraint is non-binding, all carbon taxes are set to zero; when the emissions constraint is binding, the optimal policy is characterised by the following conditions

$$\tau_{i,t}^* = \tau_t^* \ \forall \ i \in N_y$$
$$\tau_t^* = R_t \tau_1^*$$
$$\sum_{t=1}^{T} Q_{d,t} \left(\tau_t^*\right) = \bar{Q}_d$$

Proof: See Appendix B.6. Proposition 3 states that policy is intratemporally optimised when tax rates on all sectors are equalised and intertemporally optimised when the tax rate grows over time at a rate equal to the real interest rate. Both of these results are intuitive and follow from the envelope theorem, which tells us that the marginal increase in output from relaxing the emissions constraint on a given sector i at time t is equal to $\tau_{i,t}^*$. It therefore follows that to maximise the static value of GDP, a common tax rate should be applied to all sectors, and to maximise the discounted value of GDP over time, the tax rate should grow at the real interest rate, the well-known Hotelling rule.

Corollary 1: It follows from Proposition 3 that when $\tau_1^* > 0$

$$\frac{d\tau_t^*}{d\bar{Q}_d} < 0 \tag{2.45}$$
$$\frac{d\tau_t^*}{dL_t} > 0 \tag{2.46}$$
$$\frac{d\tau_t^*}{dA_t} > 0 \tag{2.47}$$
$$\frac{d\tau_1^*}{dr_{b,t}} < 0 \tag{2.48}$$

(2.45)-(2.47) state that the entire path of optimal carbon taxes increases when the carbon budget is lower, the labour force in any given period is higher or A_t in any given period is higher. In discussing these results, it is useful to note that the path of the optimal carbon tax depends on two reduced form objects: the ratio of unregulated emissions relative to the economy's carbon budget ($\frac{Q_{0d}}{\bar{Q}_d}$) and the path of the unregulated dirty energy price ($p_{0d,t}$). The larger the emissions ratio, the more the relative price of dirty to green energy has to be altered to achieve the necessary substitution away from dirty energy. Conditional on a required change in relative energy prices, a larger $p_{0d,t}$ requires a larger carbon tax. Of course, a larger $p_{0d,t}$, all else equal, would mean a substitution away from dirty energy in the laissez-faire equilibrium, lowering the emissions ratio. The emissions ratio, however, depends on the path of the dirty-to-green price ratio, so if a country were to experience a doubling of $p_{0d,t}$ and $p_{0g,t}$, for example, at all points in time, then there would be no change in $\frac{Q_{0d}}{\bar{Q}_d}$ but the carbon tax would have to double in all periods in real terms to keep emissions constant.

(2.45) and (2.46) follow from the fact that a decrease in $\bar{Q}_d$ or an increase in L_t both push up the emissions ratio. (2.47) follows from the fact that an increase in A_t (either in the form of non-energy TFP improvements, $A_{i,t}$, or neutral energy efficiency increases, $A_{ie,t}$) without any change in $A_{d,t}$ or $A_{g,t}$ has no effect on laissez-faire emissions but pushes up $p_{0d,t}$. The emissions ratio is unaffected by A_t changes since this type of TFP change does not favour the dirty or green energy sectors, leaving their laissez-faire relative price unaffected. It does, however, push up w_t, which increases $p_{0d,t}$ and requires a larger carbon tax to achieve the same proportionate increase in dirty energy prices. In the case of (2.45), (2.46) and (2.47), the entire path of carbon taxes moves in the same direction so that the necessary loosening or tightening of policy is smoothed across all periods. An increase in $r_{b,t}$ (holding constant $r_{k,t}$), lowers τ_1^*, as the policymaker finds it optimal to postpone more output losses for the future. Unlike the previous cases, however, the path of τ_t^* pivots rather than shifts, and taxes must increase at some point in the future to make up for the lower carbon taxes in the present.

2.4.2 Calibration

From equations (2.35) and (2.44) we need measures of λ_k, λ_l, A_t, $r_{k,t}$, L_t and $\{\lambda_{ie}, \tau_{i,t}, \omega_{id}, \omega_{ig}, \sigma_{ie}, A_{id,t}, A_{ig,t}\}_{i \in N_y}$ and to pin down $Q_{d,t}$.

λ_k and λ_{ie} (along with all other Cobb-Douglas expenditure share parameters of the model) are estimated using the 2016 release of the World Input-Output Tables (Timmer et al., 2015). This data set covers 43 countries and 56 sectors from 2000 to 2014. The Cobb-Douglas expenditure share parameters of the model are calibrated to match the time average of the respective expenditure shares in the data. The 43 countries contained in the data set roughly correspond to the OECD economies with the addition of nine of major developing economies. I drop eight countries[5] from this set of 43 that are large in the economic and/or climate sense[6], leaving a subset of 35 economies that conform to the 'small economy' definition used in this chapter. I also aggregate the 56 sectors contained in the original dataset to 11 sectors. This aggregation is done to facilitate the presentation of the sectorally disaggregated results in a more compact manner, but one of the primary benefits of the model is that it can readily accommodate more granularity as needed with little to no increase in computational complexity. Appendix B.1 lists all countries and sectors used.

I treat all sectors symmetrically in terms of their energy parameters such that I only need one set of parameter estimates for $\omega_d, \omega_g, \sigma_e, A_{d,t}, A_{g,t}$. This assumption means that only aggregate data on the energy mix is needed to calibrate the model, but can easily be relaxed to allow for sectoral heterogeneity in energy inputs provided sectorally disaggregated data on dirty and green energy use can be obtained. The values for $\omega_d, \omega_g, \sigma_e$ are set based on the recent study by Papageorgiou et al. (2017). As baseline values, I set $\omega_d = \omega_g = 0.5$ and $\sigma_e = 1.5$, although I conduct sensitivity analysis of my main results with $\sigma_e \in (1, 2)$.

$r_{k,t}$ and L_t are measured from the Penn World Table version 10.0, with complete time series data for all countries in my sample from 1995 to 2019. $r_{k,t}$ is taken to be the real internal rate of return in the data. I measure L_t as total hours of employment by multiplying the total number of individuals engaged with the average hours worked per person engaged. I then project these values forward for the 2020-2049 period as follows. $r_{k,t}$ is taken to be constant over this period and equal to the country-level average for the 1995 to 2019 period. To extend hours worked (L_t), I used the United Nations (2019) population projections for the growth rate of the working age population (15 to 65 years old) for each country. I use the medium variant projections as the values for the baseline calibration but also conduct sensitivity analysis using the low and high variant projections.

Historical data on $\tau_{i,t}$ is provided by the World Bank (2022) Carbon Pricing Dashboard. This dashboard provides time series data on carbon pricing across countries along with information on the percentage of emissions coverage by each carbon pricing schemes in 2012[7]. I use this data to construct measures of $\tau_{i,t}$ and the corresponding λ_{ie} for each country from 1995-2019. Since the tax price data is expressed in nominal terms, I use the Penn World Table GDP deflator to convert these values into 2017 USD PPP terms.

This leaves A_t, $A_{d,t}$ and $A_{g,t}$ to be determined. To pin down historical values for these variables from 1995-2019, I used targets for $Q_{d,t}$, $\frac{Q_{g,t}}{Q_{d,t}}$ and Y_t. I measure $Q_{d,t}$ using data

[5] Brazil, China, Germany, Indonesia, India, Japan, Russia and the United States.

[6] I exclude all countries with an average contribution to global output or emissions between 2010 and 2018 that is above 1.5%.

[7] Where multiple carbon prices are provided under one coverage data point, I take the average of all prices to be the price at which that share of emissions is covered.

on CO_2 emissions and $\frac{Q_{g,t}}{Q_{d,t}}$ from the renewable share of primary energy, both of which are compiled by Our World in Data. I obtain Y_t from the Penn World Table measure of real GDP in 2017 International Dollars. With these three targets from the data, A_t, $A_{d,t}$ and $A_{g,t}$ are then pinned down by

$$Q_{d,t} = \left(\frac{\sum\limits_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}}{\lambda_l + \sum\limits_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}} + \sum\limits_{i \in N_y} \lambda_{ie}\alpha_{ig,t}} \right) A_{d,t} L_t \tag{2.49}$$

$$\frac{Q_{g,t}}{Q_{d,t}} = \left(\frac{\sum\limits_{i \in N_y} \lambda_{ie}\alpha_{ig,t}}{\sum\limits_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}} \right) \frac{A_{g,t}}{A_{d,t}} \tag{2.50}$$

$$Y_t = \frac{w_t L_t + \sum\limits_{i \in N_y} \tau_{i,t} Q_{id,t}}{1 - \lambda_k} \tag{2.51}$$

For countries where no carbon pricing was active historically, A_t, $A_{d,t}$ and $A_{g,t}$ can be solved analytically as

$$\frac{A_{g,t}}{A_{d,t}} = \left(\frac{\omega_d}{\omega_g} \right) \left(\frac{Q_{g,t}}{Q_{d,t}} \right)^{\frac{1}{\sigma_e}} \tag{2.52}$$

$$A_{d,t} = \left(\frac{\lambda_l + \lambda_e}{\lambda_e} \right) \left[\frac{\omega_g^{\sigma_e} \left(\frac{A_{g,t}}{A_{d,t}} \right)^{\sigma_e - 1} + \omega_d^{\sigma_e}}{\omega_d^{\sigma_e}} \right] \frac{Q_{d,t}}{L_t} \tag{2.53}$$

$$A_{g,t} = A_{d,t} \left(\frac{A_{g,t}}{A_{d,t}} \right) \tag{2.54}$$

$$A_t = r_{k,t}^{\lambda_k} \left(\omega_d^{\sigma_e} A_{d,t}^{\sigma_e - 1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e - 1} \right)^{\frac{\lambda_e}{(1-\sigma_e)}} \left[(1 - \lambda_k) \frac{Y_t}{L_t} \right]^{1 - \lambda_k} \tag{2.55}$$

One major benefit of the Cobb-Douglas assumption for final demand and the sectoral production functions is that fact that the $A_{i,t}$ and $A_{ie,t}$ do not need to be separately estimated and only the composite A_t term needs to be pinned down. This keeps the calibration (and simulation) of the model tractable even when a large number of sectors are introduced in the model.

My calibrated estimates of A_t, $A_{d,t}$ and $A_{g,t}$ are then extended as follows. I assume that $A_{d,t}$ remains constant for all countries from 2020 to 2049. Between 1995 and 2019, the average annualised growth rate of my $A_{d,t}$ estimates across all countries in the sample was -0.55%, and it seems reasonable to assume that even without imposing local carbon taxes, global shifts in R&D priorities along with an increasing dependence on more expensive fossil fuel reserves will lead to stagnant $A_{d,t}$. Although I exclude the US from my analysis on account of its size, I do have data for the country and can thus go through the same calibration process to pin down the 1995 to 2019 time series for $A_{g,t}$ in the US. I then project the US estimate for $A_{g,t}$ forward at an annual rate of 2% and assume that percentage deviation in $A_{g,t}$ between each country in my sample and the US shrinks by 50% between 2019 and 2049 with $A_{g,t}$ growing at a constant rate in each country over this period. Similarly, to project A_t forward, I first project output per hour in the US forward from 2020 to 2049 by assuming it grows at an annual rate of 1.5%. I then assume that the percentage difference in output per hour between all countries in my sample and the US converges to 50% of the 2019 gap by 2049 by growing at a constant annual rate. I then back out the A_t values required to achieve this growth of output per hour in the absence of any carbon taxes using (2.55). To solve for the optimal carbon tax, I also need future projections of $r_{b,t}$. I assume a baseline value of $r_{b,t} = 2\%$ for all countries but also conduct sensitivity analysis by assuming $r_{b,t} = 0\%$ and $r_{b,t} = 4\%$. Finally, I assume a global value for the carbon budget of 550 $GtCO_2$ as of 2020. This value is consistent with restricting the global temperature increase to 1.7° with 83% probably according to IPCC estimates, as shown in Table 2.1. I also conduct sensitivity analysis letting this budget equal 300 $GtCO_2$ and 900 $GtCO_2$, consistent with meeting the 1.5° and 2.0° temperature targets, respectively, with 83% probably. These carbon budgets are then allocated to each country based on each country's share of global population in 2019 using World Bank data. The details of this calibration are summarised below in Table 2.2.

Table 2.2: Model Calibration

Parameter/Variable	Description	Value	Target/Source
ω, λ	Expenditure share parameters		WIOT
$r_{k,t}$	Marginal product of capital		PWT
$r_{b,t}$	Interest rate on borrowing		2%
L_t	Aggregate hours worked		PWT and UN
$\lambda_{ie}, \tau_{i,t}$	Carbon pricing and coverage		World Bank
ω_d, ω_g	Weight on dirty/green energy	0.50	Papageorgiou et al. (2017)
σ_e	Energy elasticity of substitution	1.5	Papageorgiou et al. (2017)
$\bar{Q}_d$	Global carbon budget	550 $GtCO_2$	IPCC
A_t, $A_{d,t}$, $A_{g,t}$	TFP		$Q_{d,t}$, $\frac{Q_{g,t}}{Q_{d,t}}$, Y_t

2.4.3 Quantitative Results

Table B.3 illustrates the optimal starting value for the carbon tax in 2020 (τ^*_{2020}) for each country under the baseline calibration assumptions. Two key points stand out from these results. First, there is significant heterogeneity in the optimal carbon tax across countries, with the baseline estimates ranging from a \$4 per ton of CO_2 in Malta to \$749 per ton of CO_2 in Taiwan. Table B.3 also provides the emissions ratio ($\frac{Q_{0d}}{Q_d}$) and the unregulated dirty energy price index in 2020 (p_{0d}) for each country. It can be seen that a high optimal carbon tax in 2020 is the result of a high emissions ratio and/or a high cost of dirty energy. In some sense, variation in the emissions ratio is the more fundamental determinant of the optimal carbon tax in that the carbon tax necessarily falls to zero once this ratio reaches one. This fact is also reflected in the data with the raw correlation between τ^*_{2020} and $\frac{Q_{0d}}{Q_d}$ being stronger than the one between τ^*_{2020} and p_{0d}, with the positive correlation between τ^*_{2020} and p_{0d} emerging when conditioned on $\frac{Q_{0d}}{Q_d}$.

In addition to the significant variation in optimal carbon taxes, a second key point is that the average optimal carbon tax value of \$325 per ton of CO_2 is significantly higher than any currently implemented or proposed carbon pricing scheme. There are three possible implications of this result. First, taking the results at face value would imply that countries are not currently doing enough to meet the climate change targets set out in the Paris Agreement. There are certainly political economy factors to suggest that climate action has been inefficiently delayed, and a wide range of research also suggests that current climate policy commitments are likely to be insufficient to limit warming to 1.7 °C (IPCC, 2021). A second possibility is that the model is missing quantitatively important aspects of reality that cause me to overestimate the level of carbon taxes required to meet the Paris Agreement. Two key areas that could be of relevance here are the assumed elasticities of substitution and the modelling of technology. In terms of elasticities of substitution, it is possible that elasticities of substitution across either final demand or intermediate inputs are larger than in my model, meaning that the required aggregate substitution away from dirty energy towards green energy happens with a much lower level of carbon taxation. Although I test the sensitivity of my results to assumptions on the elasticity of substitution between dirty and green energy, I hold all other elasticities of substitution constant and equal to one. As shown by the work of Baqaee and Farhi (2019), accounting for such changes in the production network can have large impacts on how shocks affect the aggregate economy. In terms of technology, altering the assumed growth rate of green energy productivity, expanding the set of available technologies to include abatement or negative emissions technologies, or letting technology respond endogenously to carbon taxes would also lead to a reduction in the required level of carbon taxation. A third reading of these results is that while existing carbon pricing initiatives alone will not be sufficient to meet the targets of the Paris agreement, climate policy involves a suite of policy instruments such as R&D subsidies and command and control regulation, the sum total of which may be sufficient for countries to meet their climate commitments. The reality is likely to be a mix of all three. It is also worth noting my estimates of τ^*_{2020} are well in excess of most estimates in the literature of the unconstrained optimal carbon tax or the social cost of carbon. Aside from methodological differences between studies, this difference also reflects the fact that most of these studies find an optimal value of global warming in excess of the 2.0 °C upper bound stipulated by the Paris agreement.

As well as showing the values for τ^*_{2020} by country, Table B.3 shows the loss in discounted GDP from implementing the optimal policy during the transition period (2020-2049). Given that the path of carbon taxes captures the marginal output losses from emissions regulation, it naturally follows that these output losses are closely correlated with τ^*_{2020}. The average value of these output losses across the sample is 2.1% under the baseline calibration, ranging from 0% in Malta (rounded to two decimal places) to 6.2% in Korea.

To test the sensitivity of these results to my assumptions, I assess how the average carbon tax and output loss in the sample evolves as a function of my assumptions on four key parameters/exogenous variables: the elasticity of substitution between dirty and green energy (σ_e), the real rate ($r_{b,t}$), labour supply (L_t), and the carbon budget ($\bar{Q}_d$). The results of this exercise are presented in Figure 2.4. As the carbon budget is relaxed, the real rate raised, and the elasticity of substitution across green and dirty energy is increased, the average τ^*_{2020} drops significantly. The links between $\bar{Q}_d$, r_b and τ^*_{2020} are intuitive and have been discussed in Corollary 1. The link between σ_e and τ^*_{2020} is worth discussing in a bit more detail here as it is partly dependent on the assumed path of $\frac{A_{g,t}}{A_{d,t}}$. Because I assume that $\frac{A_{g,t}}{A_{d,t}}$ is growing for all countries along the transition path, a higher σ_e decreases the emissions ratio, an effect that would be reversed if $\frac{A_{g,t}}{A_{d,t}}$ was declining. In addition, a higher σ_e always means that a given carbon tax is more effective at reducing emissions. Under the assumption that green technology improves relative to dirty technology along the transition path, both of these effects mean that raising σ_e lowers τ^*_{2020} . The average output loss associated with the optimal policy responds in a similar way as the optimal carbon tax to changes in the carbon budget and elasticity of substitution, while the average output loss is not sensitive to the real interest rate. This latter fact follows from the results discussed in Corollary 1: changes in $\bar{Q}_d$, L_t and σ_e affect the non-discounted costs of regulation whereas changes in r_b do not. Consequently, the change in r_b affects how the policymaker wishes to smooth the costs of regulation intertemporally but has little impact on the discounted sum of these costs over time. Thus, while the optimal tax value in 2020 is highly sensitive to the path of real rates over the transition period, the welfare costs of policy are not.

For the sample of countries studied, the optimal path of carbon taxation and the output cost of regulation are not sensitive to expected variation in future labour force projections. This insensitivity follows from the fact that for most of the countries in my sample, there

Figure 2.4: Average Optimal Carbon Tax in 2020

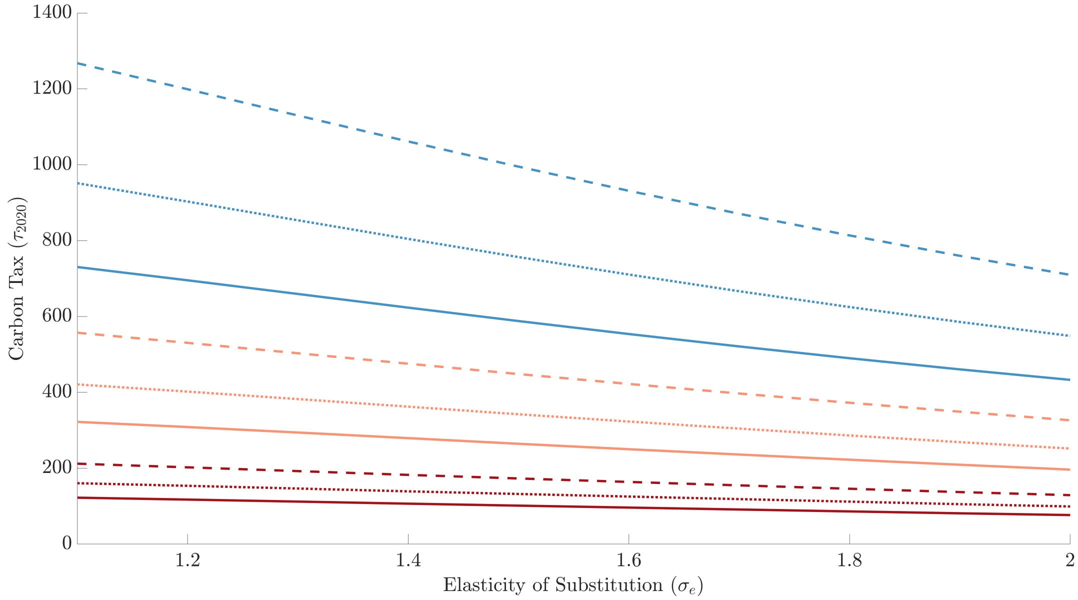

Notes: This figure illustrates the average optimal carbon tax in 2020 (τ^*_{2020}) as a function of the elasticity of substitution between dirty and green energy (σ_e). The blue, orange and red lines refer to calibrations consistent with the 1.5°C, 1.7°C and 2.0°C temperature targets, respectively; the solid, dotted and dashed lines refer to calibrations that use a value of r_b equal to 0%, 2% and 4%, respectively.

Figure 2.5: Average Discounted Value of Output Losses from Optimal Carbon Tax

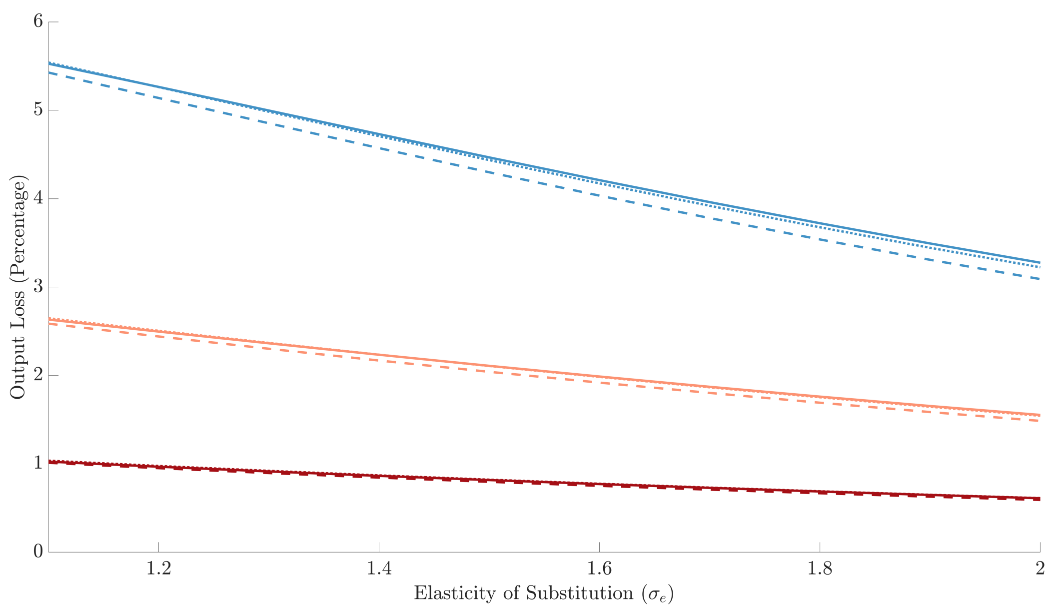

Notes: This figure illustrates the average discounted value of output losses over the transition period (as a percentage of the laissez-faire value) as a function of the elasticity of substitution between dirty and green energy (σ_e). The blue, orange and red lines refer to calibrations consistent with the 1.5°C, 1.7°C and 2.0°C temperature targets, respectively; the solid, dotted and dashed lines refer to calibrations that use a value of r_b equal to 0%, 2% and 4%, respectively.

is little variation across the main United Nations (2019) demographic projections, with the average working age population across countries expected to shrink slightly from 2020-2049. The biggest variation in the optimal carbon tax from moving from the low population projection scenario to the high one is 6.7%, while the equivalent variation in output losses is just 0.1%. Consequently, all results in Figure 2.4 are illustrated under the baseline assumption that L_t changes in line with the medium variant projection.

2.5 Suboptimal Policy: Sectoral Exemptions

This section analyses the effects of introducing sectoral exemptions to carbon pricing. The motivation for this exercise comes from the fact that most carbon pricing policies are implemented with such exemptions, which are inefficient within the context of my model. While there may be efficiency considerations missing from the model that motivate these policies, it is also possible that they are made with political economy or distributional considerations in mind. As a result, it is useful to assess the implications of such suboptimal policies when they are implemented in a benchmark environment such as the one in this chapter to understand how large the potential distortions may be.

I begin this section by characterising the suboptimal policy problem, which is to optimise the carbon tax intertemporally while providing an exogenously determined set of exemptions to specific sectors of the economy. I show that under the existing assumptions of the model and provided the tax regime is such that sectors either pay the full carbon tax or are fully exempt from the tax, there is a single dimension of suboptimality that is captured by the proportion of energy expenditures that are taxed. Reducing the dimensionality of the suboptimal policy space in this way is useful given the many ways in which sectoral exemptions can be implemented and the lack of harmonised cross-country data on these exemptions in practice. I then show the limits on the proportion of energy expenditures that can be feasibly exempt while still meeting the carbon budget, before assessing the tax and welfare implications of providing these exemptions.

2.5.1 Suboptimal Policy Problem

Given that there are an infinite range of suboptimal policies to explore, I choose to focus on one particular type of sectorally-differentiated policy: I assume that some subset of all sectors, denoted by N_0, face no carbon tax, while the remaining subset of sectors, denoted N_τ, pay a common carbon tax $\tau_t > 0$. The advantage of this assumption is that given the previous assumptions of the model, there is a unique scalar measure of the extent to which policy deviates from the efficient benchmark discussed in Section 2.4. This measure of sub-optimality turns out to be the proportion of energy expenditures that are tax exempt (see Appendix B.7 for proof), given by

$$\hat{\lambda}_{0e} \equiv \frac{\lambda_{0e}}{\lambda_e} = \frac{\sum_{i \in N_0} \lambda_{ie}}{\lambda_e}$$

Conditional on the degree of intratemporal sub-optimality, $\hat{\lambda}_{0e}$, the intertemporal policy optimisation problem is then given by

$$\max_{\{\tau_t\}_{t=1}^{T}} \quad \sum_{t=1}^{\infty} \frac{Y_t\left(\tau_t, \hat{\lambda}_{0e}\right)}{R_t}$$

$$\text{s.t.} \quad \sum_{t=1}^{T} Q_{d,t}\left(\tau_t, \hat{\lambda}_{0e}\right) \leq \bar{Q}_d$$

where τ_t is the carbon tax applied to the non-exempt sectors. Using the results established in the proof of Proposition 3 (see Appendix B.6), it follows that intertemporal optimisation still requires that the sectorally suboptimal carbon tax grows over time at a rate equal to the real interest rate.

2.5.2 Feasibility of the Emissions Constraint

When the fully optimal policy is in place, any degree of emissions reductions is feasible since aggregate emissions decline monotonically towards zero as the sectorally harmonised carbon tax approaches infinity. In contrast, with the sectorally suboptimal policy considered here, there is a lower bound on the emissions that can be attained by the carbon tax policy alone. Taking a special case of Proposition 2 with $\tau_{i,t} = \tau_t \; \forall \; i \in N_\tau$, this lower bound on emissions is given by

$$\lim_{\tau_t \to \infty} Q_{d,t} = \left(\frac{\lambda_{0e}}{\lambda_l + \lambda_e}\right) \alpha_{0d,t} A_{d,t} L_t \tag{2.56}$$

Combining (2.56) with (2.44) implies that the largest proportion of energy expenditures that can be exempt from the carbon tax while still meeting the carbon budget (which I denote by $\bar{\lambda}_{0e}$) is given by the inverse of the emissions ratio such that

$$\bar{\lambda}_{0e} = \frac{\bar{Q}_d}{Q_{0d,t}} \tag{2.57}$$

Table 2.3 shows the value of $\bar{\lambda}_{0e}$ by country under the baseline calibration assumptions and the various emissions constraints. Under the carbon budgets consistent with the 1.5 °C, 1.7 °C and 2.0 °C temperature limits, the maximum share of energy expenditures that can be feasibly exempt is 25.1%, 46.0% and 70.0%, respectively, on average across my sample. There is considerable variation around these average values, reflecting the fact that there is a large degree of variation in the extent to which each country overshoots its carbon budget in my calibrated model absent any policy intervention. In the case of the 2.0 °C emissions budget, a number of countries are able to exempt 100% of energy expenditures. This follows from the fact that these countries' emissions constraints are no longer binding in this case.

To put these numbers into context, the average share of energy expenditure in my sample accounted for by transport, manufacturing and final demand by households is 12.7%, 23.6% and 35.4%, respectively, with the other nine non-energy sectors in my model accounting for the remaining 28.3% of energy expenditures in a manner that is fairly evenly distributed. Table 2.4 shows the energy expenditures in each sector as a percentage of the maximum feasible exemptions consistent with the 1.7 °C temperature limit[8]. On average, final demand alone or manufacturing combined with transport accounts for almost all possible exemptions under the 1.7 °C carbon budget.

[8]Table B.4 in the appendix shows instead energy expenditures by sector as a percentage of total energy expenditures (that is, $\hat{\lambda}_{ie}$) across countries.

Table 2.3: Maximum Sectoral Exemption (Percentage)

	$\bar{\lambda}_{0e}$		
Country Code	1.5 °C	1.7 °C	2.0 °C
AUS	7.6	13.9	22.7
AUT	18.7	34.4	56.2
BEL	16.7	30.7	50.2
BGR	29.3	53.7	87.9
CAN	8.4	15.3	25.1
CHE	30.3	55.6	91.0
CYP	21.7	39.7	65.0
CZE	16.1	29.5	48.3
DNK	24.7	45.3	74.2
ESP	29.3	53.7	87.9
EST	16.7	30.7	50.2
FIN	16.2	29.7	48.6
FRA	28.4	52.0	85.2
GBR	23.8	43.6	71.4
GRC	26.2	48.0	78.6
HRV	35.9	65.9	100.0
HUN	34.0	62.4	100.0
IRL	15.6	28.6	46.8
ITA	27.3	50.1	81.9
KOR	13.4	24.7	40.3
LTU	39.4	72.2	100.0
LUX	7.9	14.4	23.6
LVA	44.2	81.0	100.0
MEX	42.4	77.7	100.0
MLT	54.3	99.5	100.0
NLD	16.6	30.4	49.7
NOR	14.1	25.8	42.3
POL	19.4	35.5	58.1
PRT	33.8	61.9	100.0
ROU	45.1	82.6	100.0
SVK	26.9	49.2	80.6
SVN	22.1	40.5	66.2
SWE	29.0	53.2	87.1
TUR	28.6	52.5	85.9
TWN	14.9	27.3	44.7
AVE	25.1	46.0	70.0

Notes: This table shows the maximum percentage of energy expenditures that can feasibly be exempt from carbon taxation ($\bar{\lambda}_{0e}$) under the baseline calibration assumptions and for the various carbon budgets.

Table 2.4: Energy Expenditure by Sector Relative to Maximum Feasible Exemptions

	Sector Code											
Country Code	AGR	MIN	MAN	UTL	TRD	TRN	HSP	ICT	FIN	OTH	SOC	FND
AUS	25.4	75.9	112.3	52.7	23.2	120.2	11.4	13.2	2.3	22.4	34.6	225.9
AUT	7.2	3.2	75.1	10.9	11.8	31.3	5.8	2.5	2.1	10.5	23.4	107.3
BEL	6.8	2.3	69.2	6.5	11.7	24.7	2.1	3.2	1.9	11.3	10.7	175.3
BGR	5.3	10.4	36.2	7.8	17.2	26.0	1.3	1.1	0.1	5.2	6.6	68.7
CAN	23.2	24.2	129.2	27.4	30.6	63.5	6.2	3.9	9.1	32.0	46.1	257.3
CHE	3.2	1.9	43.8	3.7	7.8	13.3	2.6	1.6	1.8	10.4	24.9	64.7
CYP	18.5	3.7	51.1	11.5	14.9	34.4	12.7	8.5	10.0	7.5	15.4	63.6
CZE	7.7	4.8	100.7	11.5	23.3	29.5	3.0	2.8	3.5	16.9	17.1	118.2
DNK	8.1	0.5	19.6	5.8	7.5	82.1	1.5	1.7	0.8	3.8	9.2	79.9
ESP	2.8	2.5	50.6	6.4	16.4	19.2	3.4	2.7	1.2	4.7	8.8	67.4
EST	20.1	3.8	59.8	17.8	16.3	71.4	3.6	2.5	1.0	14.9	24.6	90.1
FIN	8.8	4.4	79.9	11.0	10.2	44.8	2.0	2.7	2.7	42.3	20.2	107.4
FRA	5.4	0.8	40.8	6.0	12.0	18.0	1.7	5.3	1.2	7.3	10.3	83.5
GBR	4.3	1.8	50.5	6.4	12.7	13.7	3.6	4.0	4.1	8.1	13.0	107.0
GRC	10.1	0.3	19.2	4.4	7.3	16.0	2.8	1.8	0.5	3.3	6.7	135.8
HRV	4.2	3.4	23.7	8.1	14.3	12.1	5.9	4.3	2.2	9.9	14.1	49.7
HUN	6.5	1.6	51.9	4.6	9.0	17.5	1.4	1.5	0.8	5.8	7.4	52.3
IRL	6.4	13.3	72.0	16.2	28.4	14.5	3.6	21.3	11.6	38.9	32.6	91.4
ITA	4.7	1.0	48.9	9.0	16.8	18.2	3.8	1.7	1.3	6.6	13.1	74.7
KOR	5.6	3.2	150.2	13.4	14.7	49.1	4.9	3.5	2.9	15.0	26.2	116.9
LTU	6.2	0.5	11.4	2.6	2.2	9.9	0.5	0.5	0.4	2.6	5.4	96.2
LUX	7.4	0.8	217.6	16.5	27.7	112.4	10.8	12.8	80.6	64.7	39.7	101.8
LVA	11.4	0.6	22.5	6.0	8.5	25.6	1.2	1.5	0.7	8.9	9.1	27.8
MEX	2.8	2.1	27.7	4.3	7.2	25.7	3.1	0.9	0.5	5.5	7.3	41.7
MLT	2.5	0.3	19.2	3.2	7.6	21.9	4.7	0.8	1.5	6.0	2.5	30.3
NLD	11.1	6.0	59.4	7.5	13.5	33.9	3.6	1.8	1.9	9.3	13.1	168.0
NOR	12.0	10.3	54.7	10.5	13.7	42.6	2.8	2.8	1.0	11.3	29.7	196.0
POL	11.1	4.4	59.8	14.3	18.3	34.9	2.0	2.1	4.0	31.8	13.8	85.0
PRT	5.7	1.4	33.2	8.7	8.9	19.6	4.1	1.3	1.0	3.5	10.3	63.7
ROU	6.8	1.9	35.4	9.7	8.4	7.8	1.7	1.7	1.2	6.7	2.7	37.0
SVK	6.8	1.6	54.9	4.2	8.4	20.3	1.6	1.7	0.4	5.9	10.1	87.2
SVN	7.1	2.4	81.7	12.0	16.2	25.8	5.2	3.4	2.3	9.7	13.3	68.0
SWE	4.4	2.1	40.6	6.2	4.9	22.3	1.1	1.5	0.6	11.5	9.5	83.2
TUR	10.4	4.6	53.6	8.1	12.4	32.8	4.3	1.8	3.2	5.4	11.9	41.9
TWN	8.8	5.3	177.5	5.5	21.3	44.9	5.3	1.4	1.1	4.1	11.5	79.4
AVE	8.5	5.9	63.8	10.3	13.9	34.3	3.9	3.6	4.6	13.3	15.9	95.5

Notes: This table shows energy expenditure in each sector as a percentage of the total energy expenditures that can be feasibly exempt in each country under the baseline calibration assumptions.

Figure 2.6: Percentage Change in 2020 Carbon Tax due to Sectoral Exemptions

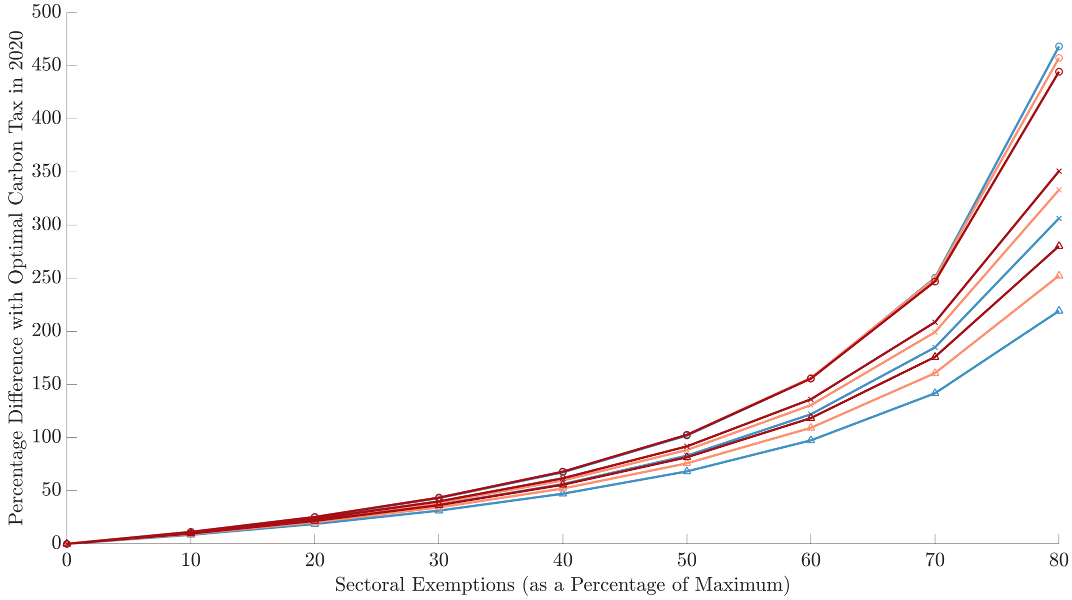

Notes: This figure illustrates the average percentage change in the 2020 carbon tax as a function of the percentage of maximum sectoral tax exemptions that are granted. The blue, orange and red lines refer to calibrations consistent with the 1.5°C, 1.7°C and 2.0°C temperature targets, respectively; the circle-, cross- and triangle-shaped markers refer to calibrations that use a value of σ_e equal to 1.1, 1.5 and 2, respectively.

Figure 2.7: Discounted Output Losses Relative to Optimal Policy (2020-2049, Percentage Difference)

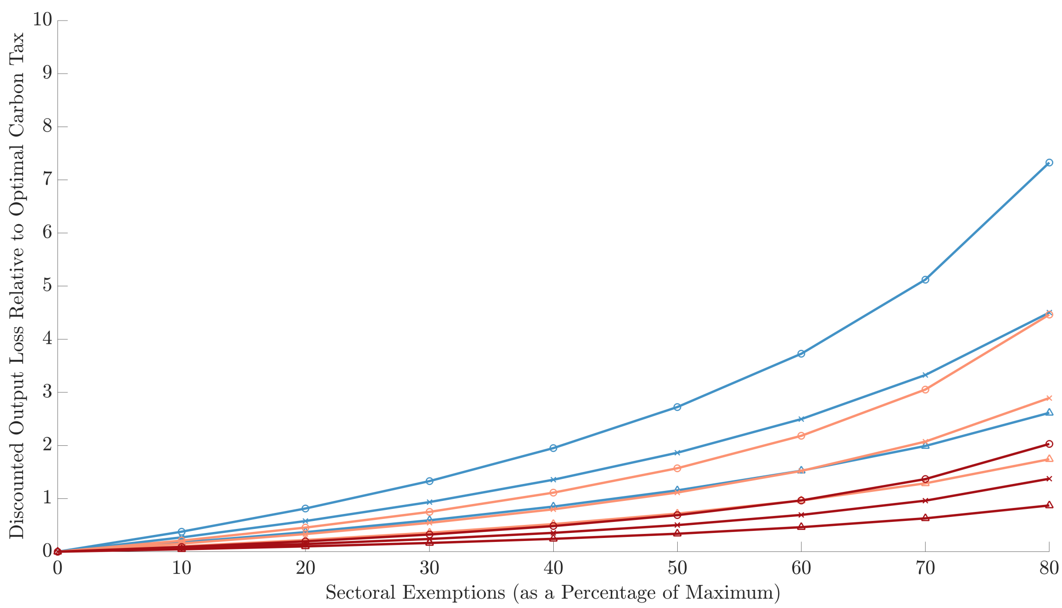

Notes: This figure illustrates the average increase in discounted output losses as a function of the percentage of maximum sectoral tax exemptions that are granted. The blue, orange and red lines refer to calibrations consistent with the 1.5°C, 1.7°C and 2.0°C temperature targets, respectively; the circle-, cross- and triangle-shaped markers refer to calibrations that use a value of σ_e equal to 1.1, 1.5 and 2, respectively.

2.5.3 Tax and Output Implications of Sectoral Exemptions

To assess the implications of providing these sectoral exemptions, I examine what happens across countries as each economy provides a given percentage of their maximum feasible exemptions $\bar{\lambda}_{0e}$. Under the baseline calibration scenario, the average path of carbon taxes increases by 59.6% when 40% of the maximum possible exemptions are granted by a country and 333.1% when 80% of the maximum possible exemptions are granted by a country. In the limit, carbon taxes approach infinity as 100% of possible exemptions are granted. The tax increase necessitated by providing such sectoral exemptions is decreasing in the elasticity of substitution between dirty and green energy as the non-exempt sectors of the economy are more responsive the higher tax rates.

Turning to the welfare implications, 0.8% and 2.9% of discounted GDP over the transition period is lost relative to the optimal policy when 40% and 80% of maximum exemptions are granted, respectively. The output losses associated with the sectorally inefficient policy are higher when the carbon budget and elasticity of substitution are lower. This follows from the fact that a lower carbon budget or elasticity of substitution requires higher carbon taxes for a given set of exemptions. This increase in carbon taxes increases the average gap between the optimal tax and the suboptimal tax across all sectors, thus increasing the distortions associated with the suboptimal tax policy. The marginal cost of these sectoral exemptions is increasing as more exemptions are provided, starting at zero and eventually becoming larger than the costs of efficient policy itself.

2.6 Expenditure Share Changes

A key assumption of the model used in this chapter is that all elasticities of substitution other than the elasticity between dirty and green energy are equal to one. This implies that expenditure shares in the model are constant with the exception of the share of energy spending going to dirty and green energy. In reality, expenditure shares are not constant, changing over the business cycle and long-term growth path. Hulten's Theorem (Hulten, 1978) tells us that a model with fixed expenditure shares can be used to approximate up to a first order the effects of shocks to a more general multi-sectoral model with intermediate input linkages. Baqaee and Farhi (2019) show that the higher order effects of shocks depend on how sales and factor shares change in response to these shocks. It follows that a key consideration in assessing whether the model presented in this chapter is likely to provide quantitatively reasonable results is the extent to which sales and factor shares are expected to change along the transition path.

Figure 2.8 shows the path of the sales and factor shares ($\lambda_{i,t}$, $\lambda_{e,t}$, $\lambda_{l,t}$, $\lambda_{k,t}$) in the data when averaged over all countries. The sales/factor shares for each country are demeaned before this averaging, showing the extent to which the data deviates from the calibrated values used in the model. An inspection of this figure suggests that most of these expenditure shares have remained relatively stable between 2000 and 2014 aside from a few notable exceptions. First, the sales share of manufacturing shows more volatility than the other sectors of the model in absolute terms. Given the large drop in the manufacturing sales share in 2009, these changes could be correlated with the business cycle but do not seem to illustrate a secular trend. Second, some sectors appear to show secular trends, most clearly in the Other Services sector but also in Financial Services and Agriculture. These changes are consistent with the broad patterns of structural transformation, with the agricultural sector shrinking and the service sector growing in relative terms.

Table B.5 shows the cross country distribution of the standard deviation for each sector/factor share from 2000 to 2014. A key point that emerges from this table is that in the Other Services and Financial Services sectors, a large share of the fluctuations over time have been driven by a handful of small countries (Ireland, Luxembourg, and Malta) whose sectoral shares changed in large part due to their special tax status and accounting-related changes in economic statistics. The secular trends in these shares in the cross-country averages should, therefore, be interpreted with caution. The volatility in the manufacturing sector observed in Figure 2.8 is, however, widespread across countries rather than being driven by a smaller number of outliers.

While it is difficult to make a precise assessment of the quantitative impact of these fluctuations in the sales/factor shares on the results of the model, the data presented here does provide some evidence to suggest that the assumptions of the model are not unreasonable. For the countries studied, the sales and factor shares are relatively constant over a 15 year period (half of the transition period in the model). The period captured by the data also included two recessions and significant energy price fluctuations. Given that most of the countries in my sample are high-income countries that are relatively mature in their process of structural transformation, assuming constant expenditure shares alongside a shift in energy spending from dirty to green energy during the transition path could be a fairly good approximation.

Although the evidence presented here is certainly not enough to be taken at face value, it does suggest that the simplified framework presented in this chapter can provide useful policy insights.

Figure 2.8: Average Sales and Factor Share Changes (Percentage Points, 2000-2014)

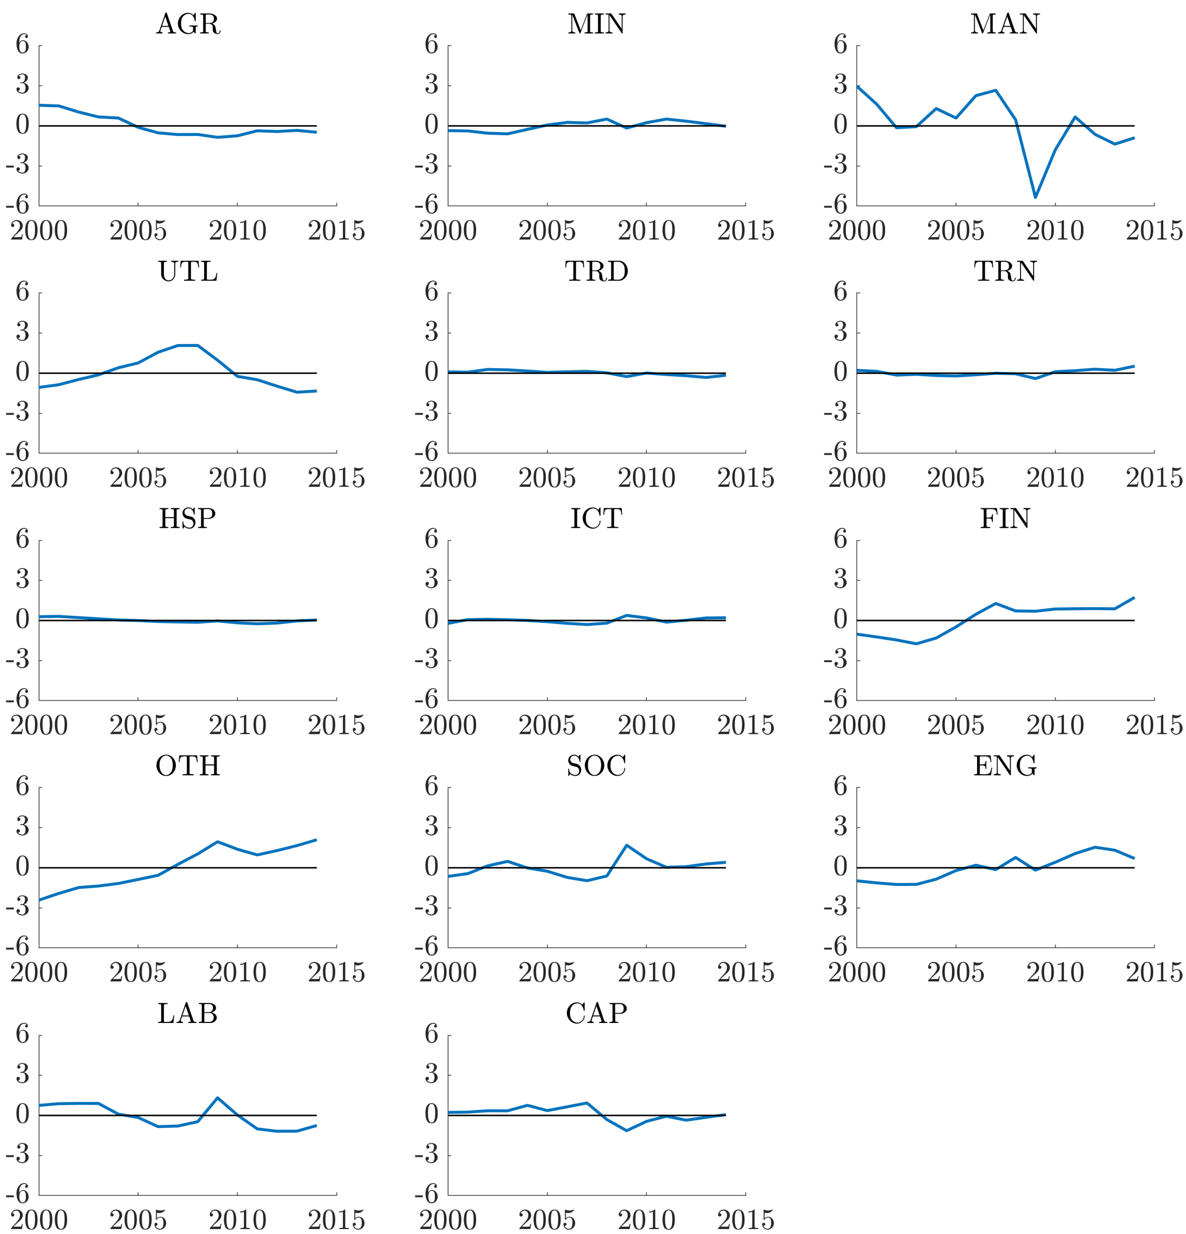

Notes: This figure illustrates path of the demeaned sales and factor shares averaged across all economies. LAB and CAP refer to the labour and capital shares, respectively (excluding value added in the energy sector).
Source: Timmer et al. (2015); Author's calculations.

2.7 Conclusion

This chapter outlines how climate policy should be formulated in small open economies that have committed to limiting themselves to a carbon budget. For the vast majority of economies, it is reasonable to take the path of climate change and interest rates on international borrowing as given. It thus follows that the optimal carbon tax is equalised across all sectors of the economy and grows in line with the real interest rate over time. There is significant variation in the optimal carbon tax across countries, with a corresponding degree of variation in the welfare costs of the policy too. While the costs of optimal policy are non-trivial, they also tend to be manageable under most assumptions, rarely exceeding 5% of the discounted value of GDP over the transition period.

By modelling a sectorally disaggregated economy, I am also able to assess the implications of providing sectoral exemptions to the carbon tax. I make the simplifying assumptions that sectors are homogenous in their use of energy inputs absent policy and either pay the full carbon tax or are completely exempt from it. In this case, the maximum share of emissions that can be exempt from the carbon tax is given by the inverse of the ratio of unregulated emissions over the transition period to the carbon budget. As a country provides a greater proportion of these maximum possible exemptions, the carbon tax required to meet its climate commitments increases as do the inefficiency costs of providing such exemptions. While the inefficiency costs of these exemptions are negligible to begin with, they eventually become significant, with the total cost of the climate policy more than doubling as the maximum feasible exemptions are granted. On average, final energy demand alone or energy demand from the manufacturing and transport sectors combined is enough to exhaust almost all the exemptions that can be feasibly granted by a country when it has a carbon budget consistent with a 1.7 °C global temperature target.

Two extensions to this paper would be particularly useful in providing further guidance to policy. First, while the structure of the model can accommodate variation in patterns of energy use and substitution across sectors, I make the simplifying assumption that all sectors are technologically homogenous in this regard. I do so to limit the data required to calibrate the model. Consequently, the only sectoral heterogeneity in energy use is the result of sectorally targeted policies. It would be interesting to allow for sectoral heterogeneity in the CES energy aggregator that is set to match sectoral data on energy use. Such a change would have potentially large implications for the results, particularly those in Sector 2.5 as there may be important interactions between policies and technology at the sectoral level.

A second useful extension would be to relax the assumption of unitary elasticities in non-energy sectors. These assumptions keep the model tractable and limit the calibration requirements, allowing the optimal policy to be easily simulated across a wide range of economies. It is also easy to understand the mechanisms driving the model, and the quantitative implications from adjusting a wide range of parameters and exogenous variables in the model can quickly be assessed. Although the evidence presented in Section 2.6 suggests that the framework used in this chapter may provide a quantitatively reasonable approximation to a more general model, it is only by extending the framework and simulating the results that this point can be formally assessed.

2.8 References

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

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.

Cai, Y. and Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy*, 127(6):2684–2734.

Cavalcanti, T., Hasna, Z., and Santos, C. (2021). Climate change mitigation policies: Aggregate and distributional effects.

Clarke, L., Jiang, K., Akimoto, K., Babiker, M., Blanford, G., Fisher-Vanden, K., Hourcade, J.-C., Krey, V., Kriegler, E., Löschel, A., et al. (2014). Assessing transformation pathways. *IPCC Assessment Report - Chapter 6*.

Devulder, A. and Lisack, N. (2020). Carbon tax in a production network: Propagation and sectoral incidence. *Banque de France Working Paper*.

Dietz, S., Gollier, C., and Kessler, L. (2018). The climate beta. *Journal of Environmental Economics and Management*, 87:258–274.

Dietz, S. and Venmans, F. (2019). Cumulative carbon emissions and economic policy: in search of general principles. *Journal of Environmental Economics and Management*, 96:108–129.

Emmerling, J., Drouet, L., van der Wijst, K.-I., Van Vuuren, D., Bosetti, V., and Tavoni, M. (2019). The role of the discount rate for emission pathways and negative emissions. *Environmental Research Letters*, 14(10):104008.

Frankovic, I. (2022). The impact of carbon pricing in a multi-region production network model and an application to climate scenarios. *Deutsche Bundesbank Discussion Paper*.

Gollier, C. (2021). The cost-efficiency carbon pricing puzzle. *CEPR Discussion Paper No. DP15919*.

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 45(3):511–518.

IPCC (2021). Climate change 2021: The physical science basis - summary for policymakers. Technical report, IPCC.

King, M., Tarbush, B., and Teytelboym, A. (2019). Targeted carbon tax reforms. *European Economic Review*, 119:526–547.

Lemoine, D. and Rudik, I. (2017). Steering the climate system: using inertia to lower the cost of policy. *American Economic Review*, 107(10):2947–57.

Nordhaus, W. D. (2008). *A question of balance*. Yale University Press.

Olijslagers, S., der Ploeg, R. v., and van Wijnbergen, S. (2021). On current and future carbon prices in a risky world.

Our World in Data (2022). Greenhouse gas emissions. *https://ourworldindata.org/greenhouse-gas-emissions Accessed on 15 April, 2022*.

Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, 99(2):281–290.

Stern, N. (2007). *The economics of climate change: the Stern review*. Cambridge University Press.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and De Vries, G. J. (2015). An illustrated user guide to the world input–output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.

Tol, R. S. (2013). Targets for global climate policy: An overview. *Journal of Economic Dynamics and Control*, 37(5):911–928.

UK Government (2022). Carbon Budgets. *https://www.gov.uk/guidance/carbon-budgets Accessed on 29 April, 2022.*

United Nations (2019). World population prospects 2019: Highlights. Technical report.

United Nations Treaty Collection (2022). Paris Agreement. *https://bit.ly/3Mn7lVS Accessed on 21 April, 2022.*

van der Ploeg, F. (2018). The safe carbon budget. *Climatic change*, 147(1):47–59.

World Bank (2022). Carbon Pricing Dashboard. *https://carbonpricingdashboard.worldbank.org/ Accessed on 3 March, 2022.*

Chapter 3

Premature Deindustrialisation and Global Integration

3.1 Introduction

The stylised facts of structural transformation are now well established: as an economy grows, the share of agriculture in the economy declines, the share of services increases, and the share of manufacturing initially increases and then declines (Herrendorf et al., 2014). A more recently documented fact is that the structural transformation process in high income countries (HICs) and low and middle income countries (LMICs) appears to be qualitatively similar but quantitatively different (Dasgupta and Singh, 2007; Felipe et al., 2019; Huneeus and Rogerson, 2020; Palma et al., 2005; Palma, 2014; Rodrik, 2016; Sposi et al., 2021). In particular, LMICs tend to have lower peak manufacturing shares (whether measured in terms of employment, or nominal or real value added) that occur at lower levels of per capita output than the early industrialising HICs, a shift referred to by Rodrik (2016) as 'premature deindustrialisation'.

The aim of this chapter is to understand the causes of this premature deindustrialisation. To do so, I proceed in three steps. First, I show that the premature deindustrialisation documented in the country-level data is less prevalent in the aggregated data for HICs and LMICs. This discrepancy between the aggregate and average industrialisation trends for these broad income groups is caused by the uneven distribution of manufacturing activity across countries, pointing to the role of international trade as a potential cause of the shifting patterns of structural transformation identified in the data. Second, I examine how the global environment faced by early and late industrialisers is likely to have differed. There has been a dramatic increase in total trade relative to GDP since the 1960s, with HICs and LMICs being similarly exposed to this unprecedented degree of global integration. In addition to deepening trade in goods, recent decades have brought new opportunities for trade in services that did not exist for early industrialisers at the beginning stages of their development. These changes in the volume and composition of international trade have also coincided with many LMICs running trade deficits that are persistent over time and large relative to the size of their economies. These deficits have been financed by the substantial aid, remittance and financial flows that are all features of the modern, globalised world. Both theory and evidence suggest that these developments could play an important role in shaping the size of the manufacturing sector in LMICs. The third and final part of the paper uses a multi-sector Eaton and Kortum (2002) model to quantitatively assess the effect of these three aspects of global integration (trade in goods, trade in services, and trade imbalances) on the sectoral composition of GDP in nine LMICs and four regional aggregates that constitute the remaining global economy economy[1].

The results of these counterfactual exercises suggest that changing patterns of global integration may be partly responsible for the observed patterns of premature deindustrialisation. Closing down trade has almost no effect on the global manufacturing share, but it does lead to a more even cross-country distribution of manufacturing across the nine LMICs I study. This is because most LMICs in the sample run trade deficits in manufacturing in the baseline equilibrium. As trade is closed, domestic production is rebalanced to match domestic demand and their manufacturing sectors expand. The two exceptions to this pattern are China and to a lesser degree India, who both ran surpluses in manufacturing in 2014 and thus see their manufacturing sectors shrink absent any trade. Trade in both goods and services play an important role in these reallocations, while the impact of trade imbalances is significantly smaller. Although the counterfactual exercises are qualitatively consistent with the notion that changing patterns of global integration have contributed to deindustrialisation pressures in LMICs, the effects tend to be small in magnitude. Across the LMICs in my model, the average gain in manufacturing share is just 1.1 percentage points, suggesting that either there are key mechanisms missing from the model and/or other driving forces of deindustrialisation are at play.

Related Literature. Although there is an extensive literature on structural transformation, much of the earlier literature focused on closed economy channels of structural transformation, particularly those arising due to the 'Baumol Effect' (non-unitary elasticities of substitution interacting with differential rates of productivity growth across sectors) and/or the 'Engel effect' (different income elasticities across sectors due to non-homothetic prefer-

[1] The nine LMICs are Bulgaria, Brazil, China, Indonesia, India, Mexico, Romania, Russia, and Turkey. See Appendix C.1 for details on the regional aggregates and sectoral aggregation scheme.

ences) (Acemoglu and Guerrieri, 2008; Boppart, 2014; Buera and Kaboski, 2009; Duarte and Restuccia, 2010; Foellmi and Zweimüller, 2008; Ngai and Pissarides, 2007). In focusing on the influence of international trade on industrialisation, I contribute to the more recently developed strand of the literature that analyses structural transformation in an open economy setting (Alessandria et al., 2021; Cravino and Sotelo, 2019; Lewis et al., 2022; Matsuyama, 2009, 2019; Sposi, 2019; Święcki, 2017; Teignier, 2018; Uy et al., 2013).

The literature around the question of premature deindustrialisation specifically is also a relatively recent development within the broader study of structural transformation. Early work by Palma (2005; 2014) and Dasgupta and Singh (2007) pointed to the existence of changing patterns of industrialisation across countries, while Rodrik (2016) coined the phrase 'premature deindustrialisation' to refer to these trends and provided broader empirical support for their existence. While there has been some push-back against the notion of premature deindustrialisation, particularly in specific regions and/or periods (Kruse et al., 2021; Nguimkeu and Zeufack, 2019), a number of subsequent studies have found that premature deindustrialisation is an empirical regularity of the cross-country data (Felipe et al., 2019; Huneeus and Rogerson, 2020; Sposi et al., 2021). I add to this evidence base by showing that much of the documented premature deindustrialisation is happening due to the cross-country distribution of manufacturing within HICs and LMICs rather than because of changing patterns of structural transformation at the global level or between HICs and LMICs as aggregate groups. The evidence I present is complementary to that of Haraguchi et al. (2017) and Sposi et al. (2021) who show that the distribution of manufacturing output across countries has become increasingly polarised over recent decades.

A few recent papers have used macro models to assess the causes of this premature deindustrialisation. Fujiwara and Matsuyama (2020) develop a model of structural transformation driven by the Baumol and Engel effects in which there are sectoral technology gaps between countries due to a lag in the diffusion of technology improvements from the frontier to follower countries. Although they do not take their model to the data, the authors show that their model can generate the premature deindustrialisation observed in the data for reasonable parameterisations. Similarly, Huneeus and Rogerson (2020) use a model of technology gaps across countries to generate cross-country differences in the process of structural transformation. Taking their model to the data, they find that they are able to account for the observed patterns of premature deindustrialisation. Aside from the exogenous diffusion of technology across countries, Huneeus and Rogerson (2020) use a closed economy model in which countries do not interact with each other. Wise (2021) extends these ideas to an open economy setting. In particular, technology gaps in his model are generated by the fact that while all countries have the same sequence of sectoral productivity growth rates over time, this common growth process is initiated in different countries at varying points in time. Due to these different initial conditions and the presence of international trade, early and late industrialisers face very different relative prices across sectors at a given point in their growth process, implying different industrialisation experiences. This chapter is distinct from these three papers in that I conduct a counterfactual exercise in which I shut down various dimensions of international trade to assess how trade has influenced the distribution of manufacturing shares across countries. By using the 'exact hat algebra' popularised by Dekle et al. (2008), I am able to conduct these counterfactual exercises without making any assumptions on the growth rates of sectoral TFP and how they vary across countries, remaining agnostic on other potential drivers of structural transformation.

Sposi et al. (2021) quantitatively assess the forces that can account for changing patterns of industrialisation across countries using a dynamic, open economy model of structural transformation. They conduct a set of counterfactual exercises that involve shutting down sectorally-biased productivity growth, shutting down international trade, and shutting down both simultaneously. They show that the interaction of the two effects is able to account for a significant proportion of the premature deindustrialisation observed in the data. My chapter is most closely related to this paper in that I also consider the effect on the cross country distribution of manufacturing shares of counterfactually shutting down international trade. My work is distinct from theirs, however, in that I decompose the effects of trade openness into four dimensions (trade in agriculture, manufacturing and services, and trade imbalances) and I calibrate my model to a wider set of LMICs. A critical shortcoming of most quantitative modelling on structural transformation is a focus on larger, higher-income countries due to data limitations. This means that few, if any, low- and lower-middle-income countries are included in this work. As I argue in Sections 3.2 and 3.3 of this chapter, understanding the structural transformation of these rarely studied economies is likely to be crucial for understanding the cross-country deindustrialisation trends observed in the data. Adding Bulgaria, Romania and Russia to my sample of LMICs makes it more representative in a number of key ways. By adding these economies to my model and decomposing the various effects of trade on industrialisation, I am able to generate additional insights that potentially shed light on the forces at play across other LMICs not in the model.

Roadmap. The rest of the chapter is organised as follows. Section 2 outlines and interprets the existing evidence on premature deindustrialisation, pointing to the key role of international trade. Section 3 provides evidence on the evolution of international trade over time that is potentially of relevance to the process of structural transformation across countries. Section 4 of the chapter outlines the quantitative model used to conduct the various counterfactuals, discusses the calibration of this model and presents the results of the quantitative exercises. Section 5 concludes.

3.2 Premature Deindustrialisation: Revisiting the Evidence

The baseline regression run by Rodrik (2016) to demonstrate the existence of premature deindustrialisation takes the form

$$m_{it} = \beta_0 + \beta_1 \ln pop_{it} + \beta_2 \left(\ln pop_{it}\right)^2 + \beta_3 \ln y_{it} + \beta_2 \left(\ln y_{it}\right)^2 + \sum_i \gamma_i D_i + \sum_T \varphi_T PER_T + \epsilon_{it} \quad (3.1)$$

where countries are indexed by i, years are indexed by t, m_{it} is the labour/value added share of manufacturing, pop_{it} is population, y_{it} is GDP per capita, D_i are country fixed effects, and PER_T are decade dummies. This regression is run using data from the Groningen Growth and Development Centre's 10-Sector Database (Timmer et al., 2015a), which provides an unbalanced panel of sectoral labour and output shares for 42 economies running from the late 1940s to the early 2010s, combined with data from the Maddison Project and the World Bank on population and GDP per capita. Rodrik (2016) finds that the estimated coefficients on the decade dummies φ_T are significant and negative, and increasing in magnitude over time (taking the 1940s and 1950s combined as the base period for the decade dummies). He also finds that these deindustrialisation trends are most pronounced for the sectoral employment share of manufacturing and least pronounced for the real value added share.

I take regression (3.1) as a starting point for my empirical analysis but rerun it using a larger cross section of countries. The United Nations (2021) provides sectorally disaggregated data on value added in real and nominal terms for almost all countries starting from 1970. I combine this data with data on total population and per capita output (measured in 2017 International Dollars) from the Penn World Tables (Feenstra et al., 2015). Combining these data sets, I construct a balanced panel of 182 countries over 50 years (1970 to 2019) with the data required to run regression (3.1). Some of the countries in this sample either split or combined during the sample period. For countries that split up (most notably, member states of the former USSR and Yugoslavia), I impute their missing data by calculating their average contribution to the parent state aggregates post independence and applying this to the historical data from the parent state. For countries that have merged, I calculate the missing historical data by aggregating the relevant data from the individual constituents. Following this procedure, the countries missing from this final panel are almost exclusively islands and city states with small populations (such as Tonga and Andorra), or fragile or

Table 3.1: Regression Results - Manufacturing Share of Nominal VA (1970-2019)

	World		HICs/LMICs		Country Level	
ln population	-119.23	(330.87)	-221.66*	(48.28)	7.59*	(0.78)
ln population squared	2.46	(7.39)	4.71*	(1.11)	0.00	(0.03)
ln GDP per capita	-32.17	(72.68)	23.44*	(7.19)	10.40*	(1.00)
ln GDP per capita squared	1.67	(3.92)	-1.19*	(0.39)	-0.50*	(0.05)
1980s	-0.85*	(0.30)	-0.58**	(0.22)	-2.03*	(0.16)
1990s	-1.78*	(0.49)	-1.45*	(0.38)	-5.52*	(0.18)
2000s	-2.99*	(0.66)	-2.09*	(0.50)	-8.48*	(0.22)
2010s	-2.36*	(0.87)	-2.10*	(0.63)	-11.12*	(0.26)
Country/Region F.E.	—		Yes		Yes	
Countries/Regions	1		2		182	
Observations	50		100		9100	

Notes: Robust standard errors are reported in parentheses. Levels of statistical significance: *99%, **95%, ***90%. World, HICs/LMICs, and Country Level refer to the results from regression 3.1 aggregated at the global level, high-income and low- and middle-income levels, and the country level, respectively.
Source: Feenstra et al. (2015); United Nations (2021); Author's calculations.

Table 3.2: Regression Results - Manufacturing Share of Real VA (1970-2019)

	World		HICs/LMICs		Country Level	
ln population	-397.62*	(123.55)	-205.86*	(40.09)	4.65*	(0.93)
ln population squared	8.85*	(2.76)	4.82*	(0.92)	-0.10*	(0.03)
ln GDP per capita	52.45***	(28.29)	-7.17	(5.41)	13.19*	(1.11)
ln GDP per capita squared	-2.66***	(1.55)	0.45	(0.30)	-0.61*	(0.06)
1980s	-0.37*	(0.10)	-0.60*	(0.17)	-0.23***	(0.12)
1990s	-0.89*	(0.18)	-1.23*	(0.23)	-1.09*	(0.14)
2000s	-0.90*	(0.24)	-1.02*	(0.30)	-2.08*	(0.16)
2010s	-0.90***	(0.46)	-1.20*	(0.40)	-3.51*	(0.19)
Country/Region F.E.	—		Yes		Yes	
Countries/Regions	1		2		182	
Observations	50		100		9100	

Notes: See Table 3.1 notes.
Source: Feenstra et al. (2015); United Nations (2021); Author's calculations.

internationally isolated LMICs (such as Afghanistan and North Korea), Taiwan being the notable exception. In total, this dataset covers over 98% of both the global population and GDP for all years in the panel.

I run regression (3.1) on this data at three levels of aggregation: for the world as a whole, for the world when split into two income groups (HICs and LMICs), and at the country-level. Figure 3.1 shows the estimated decade dummies (relative to the baseline decade of the 1970s), and Tables 3.1 and 3.2 present the full set of results from this exercise. Two key points emerge from these results. First, the country-level findings from this broader sample are consistent with those found by Rodrik (2016). A significant and progressively increasing deindustrialisation trend is evident, and this trend is larger for nominal value added than real value added. The magnitudes of the estimated effects are also broadly consistent. Second, the results show that while there is evidence to suggest that there has been downward pressure on manufacturing shares for the world as a whole since 1970, much of the premature deindustrialisation observed in the cross country data is not due to changes in the global manufacturing share or how manufacturing output is distributed between HICs and LMICs, but rather how it is distributed within these two broad income groups.

Figure 3.1: Manufacturing Value-Added Share Regression: Decade Fixed Effect Estimates

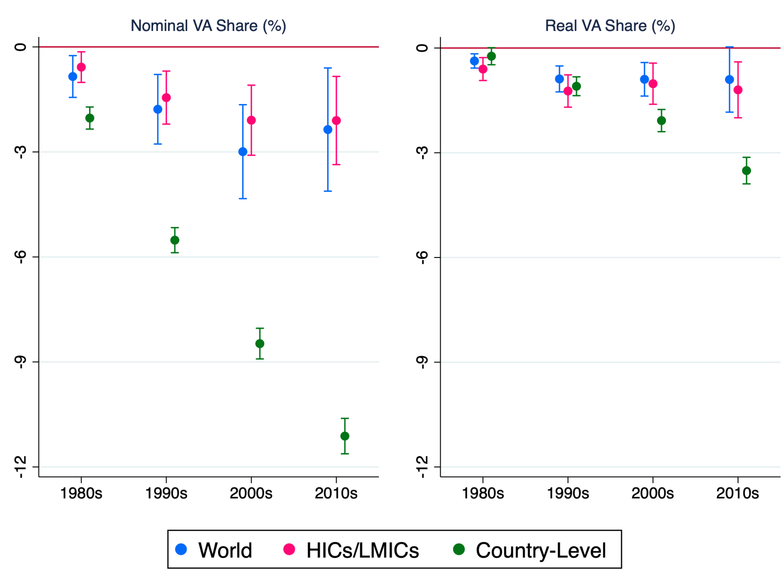

Notes: This figure illustrates the estimated decade fixed effects (relative to the baseline period of the 1970s) from running regression 3.1.
Source: Feenstra et al. (2015); United Nations (2021); Author's calculations.

The effect of averaging rather than aggregating across countries within an income grouping can clearly be seen in Figure 3.2. This figure plots the manufacturing share of output on the vertical axis and the log of GDP per capita on the horizontal axis for HICs and

LMICs. For both the nominal and real output share, I plot: (i) the aggregate manufacturing share and aggregate GDP per capita for HICs and LMICs as a whole and (ii) the average manufacturing share and average GDP per capita for countries within the HIC and LMIC groupings[2]. When looking at the figures produced by the income group averages (which reflect the same level of aggregation as the country-level regressions), the case for premature deindustrialisation amongst LMICs is evident. In terms of the average value added share in nominal terms, the LMIC curve is shifted down and to the left in comparison to the HIC curve. As with the regression results, the average value added share in real terms appears to have held up more in LMICs against this premature deindustrialisation trend. While the average peak manufacturing share may be similar across the two groups in real terms, however, the timing of this peak still appears to occur at much lower levels of income per capita for LMICs.

In contrast, the scatter plots for the aggregate manufacturing share show a very different picture. Although LMICs as a group are deindustrialising in nominal terms and this deindustrialisation appears to have set in at a lower level of per capita income than for HICs as a whole, there is no clear evidence to suggest they experienced a lower peak share of manufacturing or that this share will be lower than the HIC share once they reach comparable levels of per capita GDP. The change is even more dramatic when looking at the aggregated data for the real value added share of manufacturing: LMICs as a group experienced five decades of significant industrialisation from 1970 to 2019 and it is not clear that they have yet reached their deindustrialisation phase. The current real share of manufacturing is also significantly higher than the real share for the HIC aggregate at any period in the data, although it is possible that the real share for HICs as an aggregate peaked before the start period of the sample (1970).

Taken together, the evidence presented in this section suggests that premature deindustrialisation is not so much a phenomenon operating at the aggregate level of the global economy or between LMICs and HICs, although there certainly may be contributing factors operating at this level, but is rather a phenomenon that has been driven by changes in the cross country distribution of manufacturing output within LMICs themselves. This reading of the data is consistent with the findings of Haraguchi et al. (2017) and Sposi et al. (2021) who both find that the polarisation of manufacturing activity has increased in recent dec-

[2] While Haraguchi et al. (2017) conduct a similar exercise, they present the time series of the relevant manufacturing shares for LMICs and HICs. While it is somewhat instructive to see how the discrepancies between the average and aggregate shares have evolved over time for LMICs and HICs, the key question from the perspective of the structural transformation literature is whether the relationship between the manufacturing share and GDP per capita is systematically different for HICs and LMICs. It is, therefore, essential to condition the results relating to manufacturing share changes on the relevant measure of GDP per capita, as I do in Figure 3.2.

Figure 3.2: Manufacturing Share of Value Added (1970-2019)

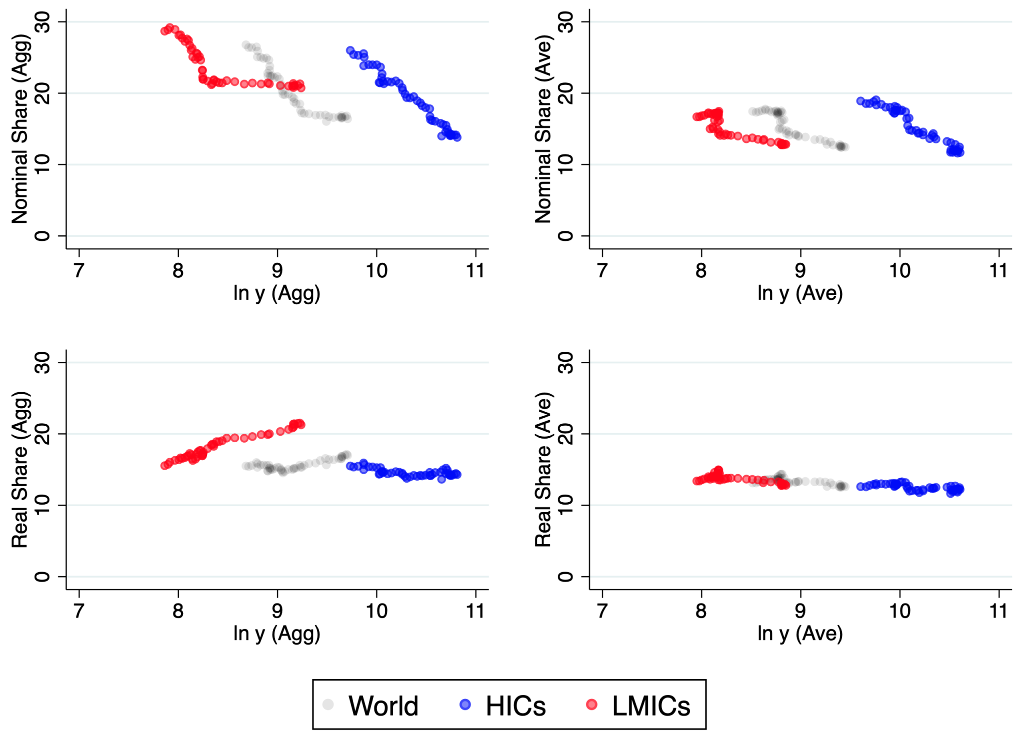

Notes: This figure illustrates the manufacturing share of value added against the log of income per capita (2017 PPP USD) for high income countries (HICs), low and middle income countries (LMICs) and the world as a whole. Subplots in the first column show the results when aggregating across all countries in the respective groups; subplots in the second column show the results when averaging across all countries in the respective groups.
Source: Feenstra et al. (2015); United Nations (2021); Author's calculations.

ades, with manufacturing output increasingly being concentrated in a relatively small set of countries, many of which tend to have above average population sizes. It therefore follows that any attempt to explain the premature deindustrialisation observed in the data should take into consideration the role of international trade.

3.3 Premature Deindustrialisation: The Role of International Trade

Broadly speaking, international trade may have contributed to premature deindustrialisation in two ways. First, the trading partners of late industrialisers have different characteristics than the trading partners of early industrialisers at comparative stages of their development. By importing lower relative prices for manufactures through either technology diffusion or import competition (the Baumol effect) combined with higher demand for services from higher income countries conditional on relative prices (the Engel effect), late industrialisers face deindustrialisation pressures sooner than earlier industrialisers. These are the channels emphasised in much of the existing work that attempts to account for premature deindustrialisation in LMICs (Fujiwara and Matsuyama, 2020; Huneeus and Rogerson, 2020; Wise, 2021). Second, conditional on a given set of trading partners and their characteristics, late industrialisers are integrated with their trading partners in a very different manner than earlier industrialisers were at the begging of their industrialisation process. The evidence presented in Section 3.2 suggests that these changing patterns of global integration have played an important role given that much of the premature deindustrialisation observed in the data appears to be driven by the distribution of manufacturing activity across LMICs rather than an aggregate loss of manufacturing activity in LMICs as a whole. To motivate the counterfactual exercises I conduct in Section 3.4 of this chapter, I outline three key dimensions of global integration that are likely to differ from the early period of industrialisation in HICs.

3.3.1 Total Trade

In terms of the overall volume of trade, the typical LMIC is likely to be more globally integrated than the typical HIC at a given stage in their industrialisation process. The total volume of international trade (exports plus imports) relative to world GDP has increased to historically unprecedented levels since the 1960s. The so-called first era of globalisation lasted from the mid 1800s to the beginning of the first World War. As illustrated in Figure 3.3, global trade as a percentage of GDP peaked at just under 30% before the First World War. While this peak was substantially higher than estimates of global trade openness in preceding periods, it is still small relative to the contemporary era of globalisation that started in the 1960s. The current wave of globalisation has lead to a doubling of trade flows relative to the size of the global economy, with trade to GDP peaking at 61% before the Great Recession in 2008.

Importantly, this significant increase in global integration has been experienced across all income groups in the cross-country income distribution. As shown by 3.4, although HICs have tended to have a slightly higher degree of openness than LMICs in the aggregate, any dispersion in trade openness across income groups has been relatively small (especially from 1990 up to the 2008 recession). Given the historical measures of trade openness in Figure 3.3 at the global level, it is likely that the typical LMIC has been much more exposed to global trade than the typical HIC at equivalent points in their respective industrialisation.

Figure 3.3: Total Trade as a Percentage of GDP (World, 1870-2017)

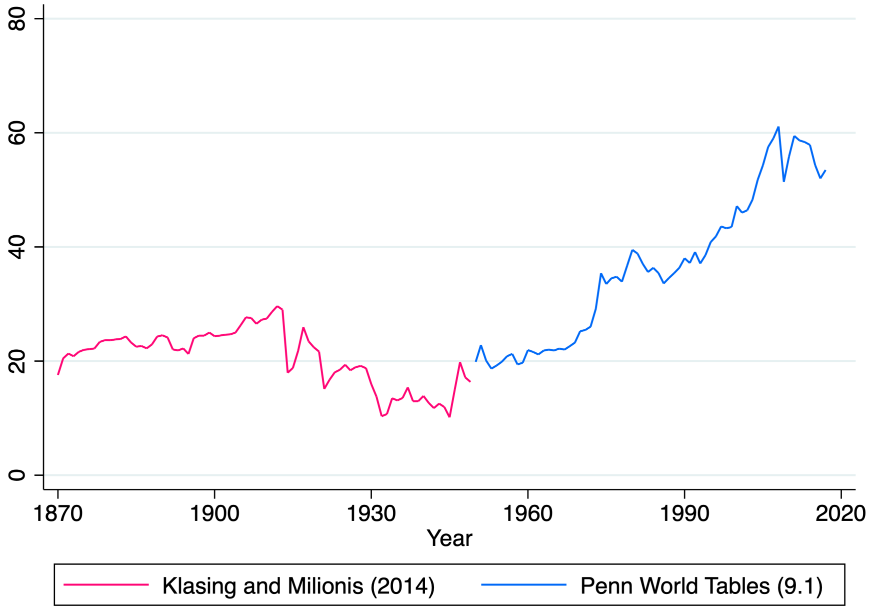

Notes: This figure illustrates the sum of world exports and imports as a percentage of world GDP. The estimates from 1870 to 1949 come from Klasing and Milionis (2014); the estimates from 1950 to 2017 come from Feenstra et al. (2015).
Source: Our World in Data (2022)

Figure 3.4: Total Trade as a Percentage of GDP (1960-2021)

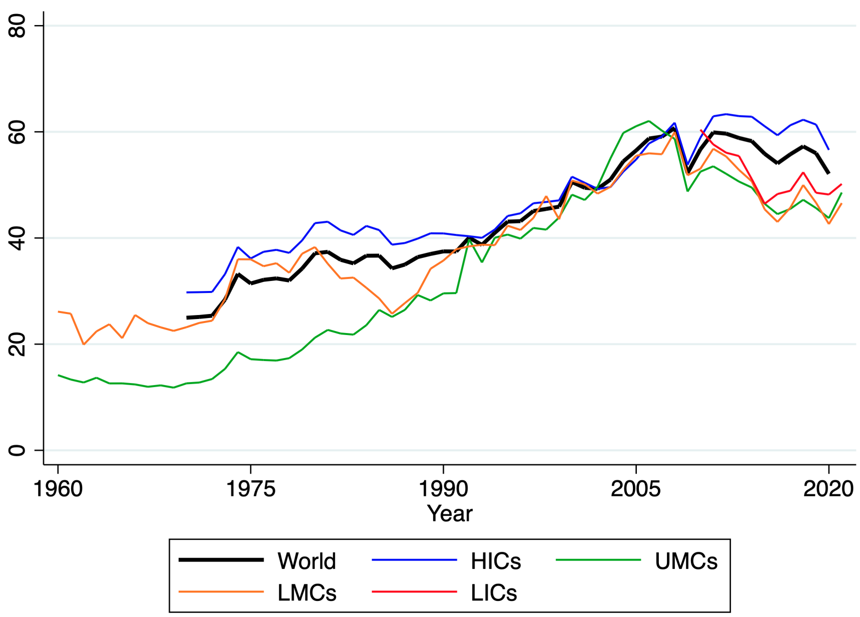

Notes: This figure illustrates the sum of exports and imports as a percentage of GDP by income group. Income groups (as classified by the World Bank) are abbreviated as follows: High Income Countries (HICs), Upper-Middle Income Countries (UMCs), Lower-Middle Income Countries (LMCs), Low Income Countries (LICs).
Source: World Bank (2022)

3.3.2 Trade in Services

In addition to a deepening of trade, there have also been significant changes in the composition of global trade over the last fifty years. One notable development has been the rise of the trade in services. Figure 3.5 shows the large increase in global trade in services from 1975 to 2020. While HICs have been more exposed than average to this trend and there is more substantial variation in the exposure to trade in services than trade in goods across the broad income categories, LMICs have been meaningful participants in such trade.

Figure 3.5: Total Trade in Services as a Percentage of GDP (1970-2021)

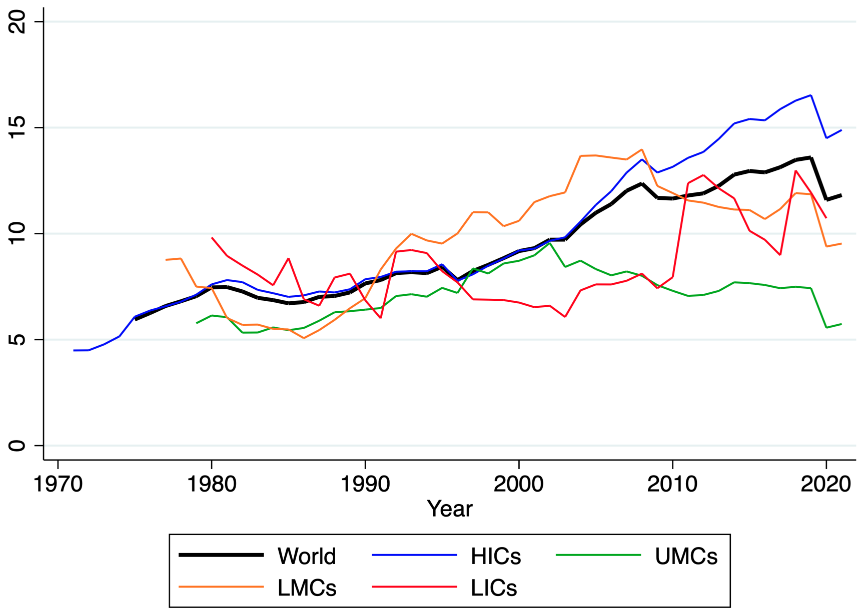

Notes: This figure illustrates the sum of exports and imports of services as a percentage of GDP by income group. Income groups (as classified by the World Bank) are abbreviated as follows: High Income Countries (HICs), Upper-Middle Income Countries (UMCs), Lower-Middle Income Countries (LMCs), Low Income Countries (LICs).
Source: World Bank (2022)

Many services have historically been non-tradable as they required the producer and consumer to be physically present in the same location for the service to be delivered. Technological changes in recent decades, however, have facilitated the increase in the international trade in services in two fundamental ways. First, technological development in the telecommunications sector has meant that many services that previously required physical proximity can now be delivered remotely. Figure 3.6 shows the growth of services defined as 'digitally deliverable' by UNCTAD[3]. These services form a substantial and growing share of total service exports across all income groups (although some LMICs are still significantly less reliant

[3]Digitally-deliverable services are defined by UNCTAD STAT as 'an aggregation of insurance and pension

on such exports than HICs). While the data reflect the share of service exports that could be digitally traded ('digitally deliverable') rather than those that are actually digitally traded ('digitally delivered'), these trends are still instructive. The data show that conditional on the required ICT infrastructure being in place, many LMICs have the necessary comparative advantage to export digital services to the world. Given the increasing penetration of ICT infrastructure in LMICs, it stands to reason that the rise of 'digitally deliverable' services has been accompanied by a rise in 'digitally delivered' service exports, and this trend is likely to continue moving forward. To the extent that this is the case, the information technology revolution has opened up opportunities for service exports that would have been unavailable to the early industrialising HICs.

Figure 3.6: Percentage of Service Exports that are Digitally-Deliverable (2005-2020)

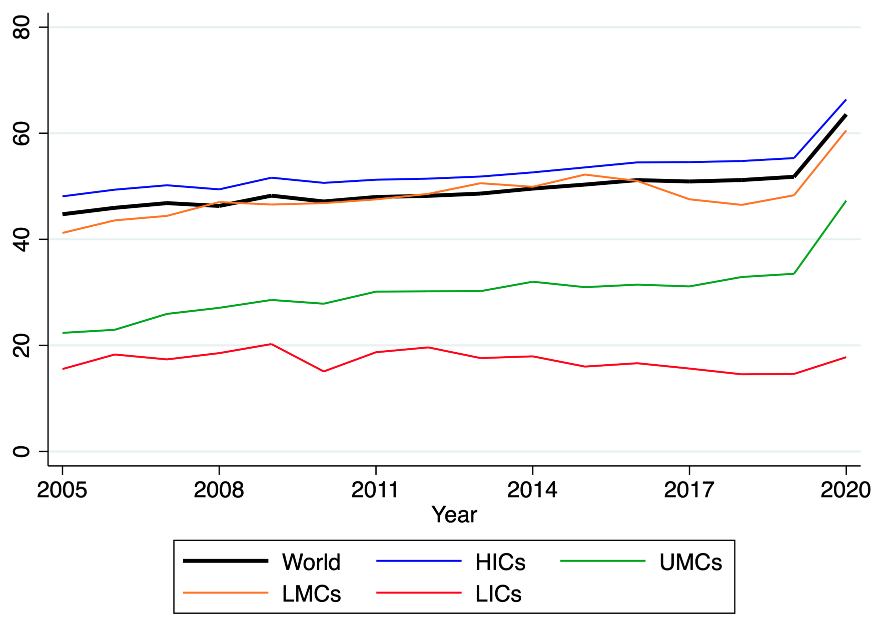

Notes: This figure illustrates the percentage of service exports that are classified as digitally-deliverable by income group. Income groups (as classified by the World Bank) are abbreviated as follows: High Income Countries (HICs), Upper-Middle Income Countries (UMCs), Lower-Middle Income Countries (LMCs), Low Income Countries (LICs). Digitally-deliverable services are defined by UNCTAD STAT as 'an aggregation of insurance and pension services, financial services, charges for the use of intellectual property, telecommunications, computer and information services, other business services and audiovisual and related services'. Source: UNCTAD STAT (2022)

A second fundamental change to the non-tradability of services has followed from advances in the transportation sector, particularly the aviation industry, that have lead to the development of a mass market for international travel. As shown by Figure 3.7, flights taken

services, financial services, charges for the use of intellectual property, telecommunications, computer and information services, other business services and audiovisual and related services'.

Figure 3.7: Ratio of Air Passengers Carried to World Population (1970-2020)

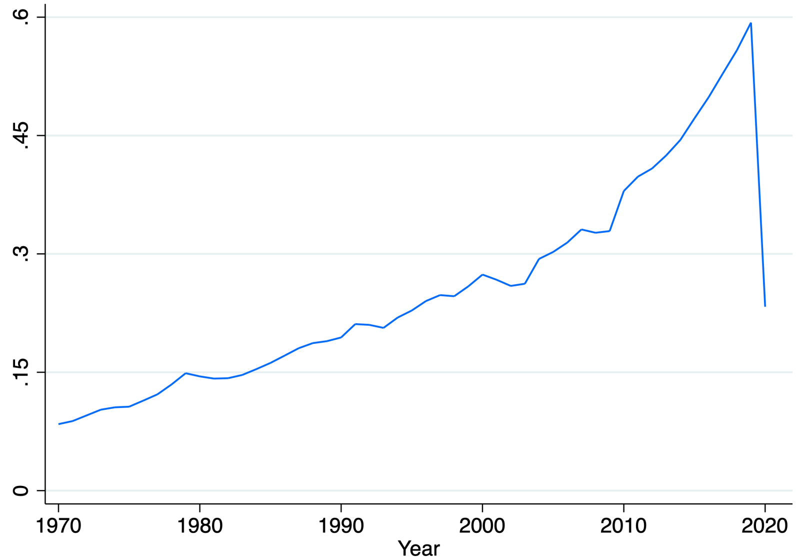

Notes: This figure illustrates the total number of air passengers carried globally (both domestic and international) divided by the total world population on an annual basis from 1970 to 2020.
Source: World Bank (2022); Author's calculations.

Figure 3.8: Tourism as a Percentage of Total Exports and ln GDP Per Capita by Country

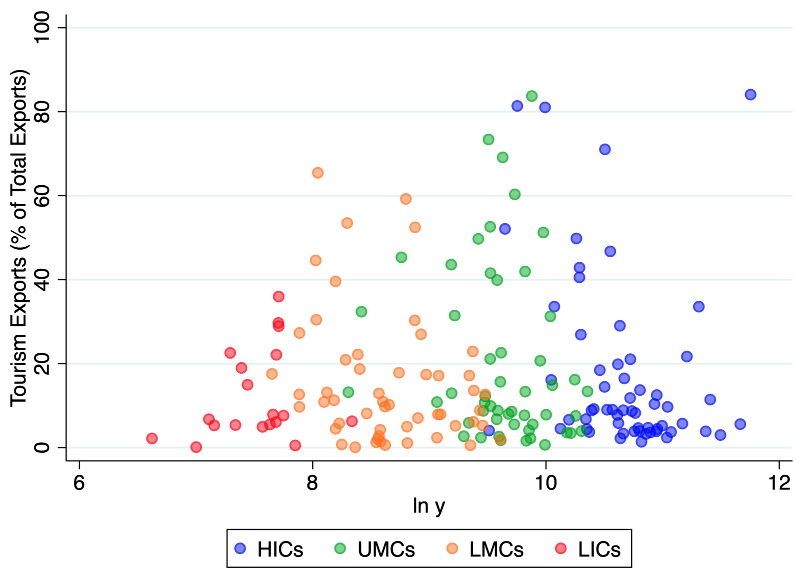

Notes: This figure illustrates average international tourism receipts as a percentage of total exports from 1995 to 2019 and the natural log of GDP per capita in 2019 (measured in 2017 PPP USD) by country. All countries with at least one data point for tourism receipts between 1995 and 2019 are included (a total of 182 countries). Of these countries, 169 have data for all years between 1995 and 2019. See Figure C.3 in Appendix C.3 for details on the missing data.
Source: World Bank (2022); Author's calculations.

per capita at the global level increased over sevenfold from 1970 to 2019 (the sharp decline in 2020 representing the Covid shock). As discussed with other types of trade, the impact of this rise of mass international travel has been felt across all income groups. Figure 3.8 shows average international tourism receipts as a percentage of exports between 1995 and 2019 by country as a function of each country's log GDP per capita in 2019. It is clear from this figure that there are a large number a LMICs who rely heavily on international tourism for export revenue. Although no comparable data exists for the early era of HIC industrialisation, it seems safe to assume that only a handful of economies, if any, would have been similarly dependent on tourism as an export during this era given the fairly recent emergence of mass tourism.

This rise in the tradability of services during the current era of globalisation has provided opportunities for late industrialisers to exploit potential comparative advantage in service sectors that would have not been available to early industrialisers. While it is still true that export growth is a necessary condition of development for all but possibly the largest of economies, the data presented above shows that the composition of exports can now be much more tilted towards the service sector. This means that some LMICs can potentially grow, diversify and integrate with the global economy by shifting resources directly from agriculture to services without ever developing a large manufacturing base.

3.3.3 Trade Imbalances

A final dimension of the modern era of globalisation worth mentioning is the existence of large and persistent trade imbalances across countries. Figure 3.9 shows the cross-country distribution of average trade balances as a percentage of GDP between 1995 and 2019. The median value of this distribution is -3.7% of GDP, with the average balance only reaching 0% of GDP at the 65th percentile of the distribution. Amongst countries whose average balance was negative, the cross-country average of the average balance was -14.2% of GDP; the corresponding figure for countries whose average balance was positive was 9.7% of GDP. In addition to the distribution of average trade balances relative to GDP being skewed in these ways, deficit and surplus countries tended to have different characteristics. In particular, countries that ran deficits on average also tended to be far smaller and poorer than surplus countries, with an average population size of 23.9 million and average GDP per capita of $7,624 (in 2017 International Dollars) compared to average values of 38.9 million and $27,138 for surplus countries, respectively. In sum, a relatively small number of larger, higher income countries have persistently run trade surpluses that support the persistent trade deficits of a relatively larger group of smaller, lower income countries, with the trade imbalances in the surplus countries tending to be smaller relative to the size of their economies than the imbalances in the deficit countries. I further verify this result by running the following regression

$$Balance_{it} = \beta_0 + \beta_1 \ln y_{it} + \beta_2 \ln Y_{it} + \beta_3 Oil_{it} + \sum_{t=1996}^{2019} \varphi_t Year_t + \epsilon_{it} \quad (3.2)$$

where i indexes countries, t indexes years, $Balance_{it}$ is country's trade balance as a percentage of their GDP in year t, y_{it} is GDP per capita, Y_{it} is aggregate GDP, Oil_{it} is a country's oil rents as a percentage of their GDP, and $Year_t$ is a year fixed effect. The results from running regression 3.2 show that within a given year, poorer and smaller countries tend to run larger deficits relative to GDP.

Figure 3.9: Cross-Country Distribution of Trade Balances (% of GDP, 1995-2019 Average)

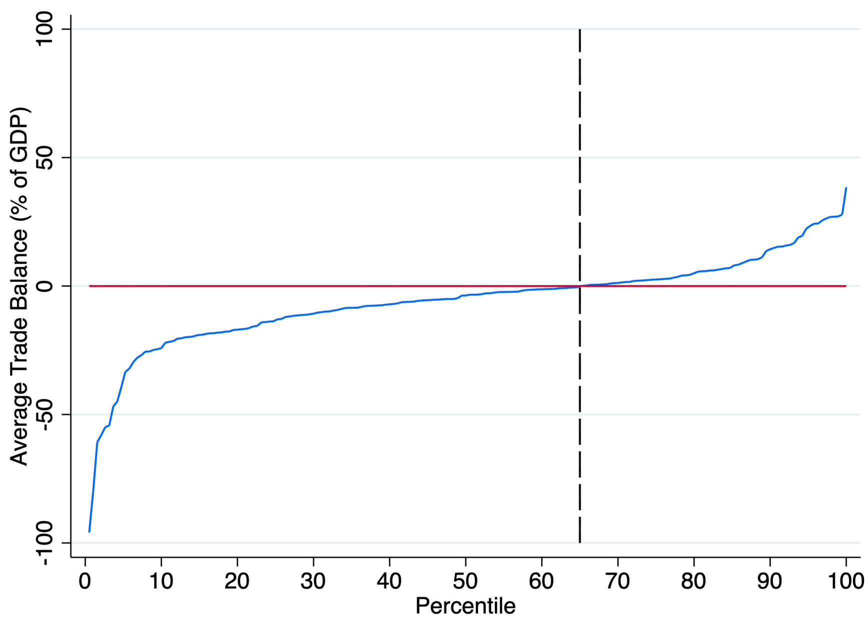

Notes: This figure illustrates the cross-country distribution of each country's average trade balance as a percentage of their GDP from 1995 to 2019. All countries with at least one data point for the relevant period are included (a total of 190 countries). Of these countries, 157 have data for all years between 1995 and 2019. See Figure C.1 in Appendix C.2 for details on the missing data.
Source: World Bank (2022); Author's calculations.

Table 3.3: Regression Results - Trade Balance as a Share of GDP (1995-2019)

	Unbalanced Panel				Balanced Panel			
ln GDP per capita	1.88*	(0.14)	2.04*	(0.15)	1.33*	(0.13)	1.47*	(0.13)
ln GDP	5.28*	(0.21)	5.86*	(0.22)	4.69*	(0.21)	5.07*	(0.22)
Oil Rents (% of GDP)	0.60*	(0.02)	—	—	0.56*	(0.03)	—	—
Year Fixed Effects	Yes		Yes		Yes		Yes	
Country Fixed Effects	No		No		No		No	
Countries/Regions	177		177		145		145	
Observations	4,114		4,114		3,625		3,625	

Notes: Robust standard errors are reported in parentheses. Levels of statistical significance: *99%, **95%, ***90%. See Figure C.2 in Appendix C.2 for details on data discrepancies between the unbalanced and balanced panel.
Source: World Bank (2022); Author's calculations.

LMICs have financed these persistent deficits using a mix of aid, remittances, and financial inflows, all of which have grown substantially during the current period of globalisation. Trade deficits are relevant to the manufacturing share of the economy because they allow for a lower level of tradable output for a given level of tradable consumption. To the extent that the agriculture and manufacturing sectors are more tradable than the services sector (which the data suggests is still the case, despite the rise in trade in services mentioned previously[4]), this would imply that growing trade deficits could be associated with the simultaneous shrinking of agriculture and manufacturing alongside the rise of services in low income countries. This type of deficit driven deindustrialisation has theoretical and empirical support from of the Dutch Disease literature (Acosta et al., 2009; Corden and Neary, 1982; Lartey et al., 2012; Rajan and Subramanian, 2011), and the sectoral implications are also consistent with much of the evidence on premature deindustrialisation (Gollin et al., 2016). Given that the trade deficit countries tend to be smaller, the proportionate impact of these deficits on their sectoral composition is likely to be larger than the impact on the surplus countries implying that this could lead to a shift in the average cross country manufacturing share. While it is difficult to obtain data on trade balances and GDP for a wide range of economies during the early era of industrialisation, it seems unlikely that the early industrialisers would have sustained such large and persistent deficits when they were industrialising.

Taken together, the evidence presented in this section suggests that there are three key ways in which global integration has differed during early and late stage industrialisation: trade linkages have deepened, the composition of exports has tilted towards the service sector

[4]See, for example, Lewis et al. (2022).

for a substantial number of countries, and many countries now run large and persistent trade deficits. The next section of the paper evaluates quantitatively the effect of these changes on sectoral value added shares across countries.

3.4 Counterfactual Results

In this section of the chapter, I use a multi-sector version of the Eaton and Kortum (2002) trade model developed by Caliendo and Parro (2015) to conduct a number of counterfactual exercises that shut down various types of integration in the global economy. The first part of the section presents the model and the equilibrium conditions. I then discuss how I calibrate the model and take it to the data, before moving on to the results from the counterfactual exercises that I run.

3.4.1 Model Setup

The world economy is made up of N countries that produce and consume goods and services in J sectors. Each country n has a representative household with preferences given by

$$u\left(C_n\right) = \prod_{j=1}^{J} C_n^{j\,\alpha_n^j} \qquad \sum_{j=1}^{J} \alpha_n^j = 1 \tag{3.3}$$

where C_n^j is consumption in country n of final output from sector j. Household income (I_n) in country n comes from two sources: households supply labour (L_n) for which they receive wage income (w_n), and they receive lump sum income transfers across countries in the form of trade deficits (D_n). Both L_n and D_n are exogenously determined in the model.

A continuum of differentiated intermediate goods $\omega^j \in [0,1]$ is produced in each sector j using labour and composite intermediate goods as inputs. The production function for intermediate good ω^j in sector j in country n is given by

$$q_n^j\left(\omega^j\right) = z_n^j\left(\omega^j\right)\left[l_n^j\left(\omega^j\right)\right]^{\gamma_n^j} \prod_{k=1}^{J}\left[m_n^{k,j}\left(\omega^j\right)\right]^{\gamma_n^{k,j}} \qquad \sum_{k=1}^{J} \gamma_n^{k,j} = 1 - \gamma_n^j \tag{3.4}$$

where $z_n^j(\omega^j)$ is TFP, $l_n^j(\omega^j)$ is labour input, and $m_n^{k,j}(\omega^j)$ is a composite intermediate input from sector k. The productivity term $z_n^j(\omega^j)$ is drawn from a Fréchet distribution with a location parameter λ_n^j and a shape parameter θ^j. The distribution of productivities is assumed to be independent across goods, sectors and countries. Production is constant returns to scale and markets are perfectly competitive, implying that firms price at unit cost. It follows that prices can be expressed as

$$p_n^j\left(\omega^j\right) = \frac{c_n^j}{z_n^j\left(\omega^j\right)}$$

where c_n^j is the unit cost component common to all firms operating in sector j in country n. The expression for c_n^j is given by

$$c_n^j = \left(\frac{w_n}{\gamma_n^j}\right)^{\gamma_n^j} \prod_{k=1}^{J} \left(\frac{P_n^k}{\gamma_n^{k,j}}\right)^{\gamma_n^{k,j}}$$

where P_n^k is the price of the composite intermediate good from sector k in country n.

The supply of the composite good from sector j in country n is given by Q_n^j. This good is created at minimum cost by purchasing intermediate goods ω^j from the lowest cost supplier globally. The production technology for the composite is a CES aggregator:

$$Q_n^j = \left[\int r_n^j(\omega^j)^{1-1/\sigma^j} d\omega^j\right]^{\sigma^j/(\sigma^j-1)} \tag{3.5}$$

The composite Q_n^j is sold to consumers and intermediate producers in country n at unit cost. It is assumed that the elasticity of substitution (σ^j) and the Fréchet shape parameter of the productivity distribution (θ^j) are such that $1+\theta^j > \sigma^j$. Trade costs take the form of standard iceberg trade costs, whereby $d_{ni}^j > 1$ units of the intermediate good in sector j need to be produced in country i for one unit of the good to be delivered to country n. It is assumed that $d_{ni}^j = 1$ and that $d_{ni}^j < d_{nm}^j d_{mi}^j$ for any three countries n, m and i.

Total expenditure on sector j goods in country n is denoted by $X_n^j = P_n^j Q_n^j$. Expenditure in country n on sector j goods imported from country i is given by X_{ni}^j. It follows that the share of sector j expenditure in country n that is spent on goods originating in country i is given by

$$\pi_{ni}^j = \frac{X_{ni}^j}{X_n^j}$$

I proceed by solving for the equilibrium using the 'exact hat algebra' used by Dekle et al. (2008). Eaton and Kortum (2002), Dekle et al. (2008), and Caliendo and Parro (2015) provide detailed explanations of the derivation and solution of these conditions, so I provide just a brief summary here. For any variable x in a baseline equilibrium, let x' denote the value of that variable in a counterfactual equilibrium and $\hat{x} = \frac{x'}{x}$ be the relative change in the variable induced by moving from the baseline equilibrium to the counterfactual one. I use the 2014 values for GDP and trade flows observed in the data to construct the baseline equilibrium of the model. I then consider five counterfactual changes to this baseline: separately eliminating trade in agriculture, manufacturing and services; eliminating trade imbalances; and imposing autarky on all countries. Since these counterfactual changes are implemented via exogenous shocks to trade costs and trade deficits, the equilibrium system of equations in relative changes in all cases is given by:

$$\hat{c}_n^j = \hat{w}_n^{\gamma_n^j} \prod_{k=1}^{J} \left(\hat{P}_n^k\right)^{\gamma_n^{k,j}} \tag{3.6}$$

$$\hat{P}_n^j = \left[\sum_{i=1}^{N} \pi_{ni}^j \left[\hat{d}_{ni}^j \hat{c}_i^j\right]^{-\theta^j}\right]^{-1/\theta^j} \tag{3.7}$$

$$\hat{\pi}_{ni}^j = \left[\frac{\hat{d}_{ni}^j \hat{c}_i^j}{\hat{P}_n^j}\right]^{-\theta^j} \tag{3.8}$$

$$X_n^{j\prime} = \sum_{k=1}^{J} \gamma_n^{k,j} \sum_{i=1}^{N} \pi_{in}^{k\prime} X_i^{k\prime} + \alpha_n^j I_n^\prime \tag{3.9}$$

$$\sum_{j=1}^{J} \sum_{i=1}^{N} \pi_{ni}^{j\prime} X_n^{j\prime} - D_n^\prime = \sum_{j=1}^{J} \sum_{i=1}^{N} \pi_{in}^{j\prime} X_i^{j\prime} \tag{3.10}$$

$$I_n^\prime = \hat{w}_n w_n L_n + D_n^\prime \tag{3.11}$$

Condition (3.6) is the change in unit costs in sector j of country n induced by the relevant counterfactual. Since fundamental productivities are left constant, this change is common to all intermediate variety producers in a given sector-country pair. (3.7) is the change in the sectoral price index, and (3.8) is the change in bilateral trade shares. (3.9) is the expression for total expenditure in country n on sector j goods in the counterfactual equilibrium, while (3.10) is the trade balance equation (adjusting for the new trade deficit) in the counterfactual equilibrium. Finally, (3.11) is an expression for total household income under the counterfactual.

3.4.2 Model Calibration

The benefit of adopting this approach is that I can analyse the impact of various counterfactual changes to global integration on the manufacturing share across countries without having to estimate a number of parameters and exogenous variables (notably, productivities and trade costs). To solve the equilibrium system of equations (3.6)-(3.11), I require data on baseline trade flows (π_{ni}^j) and GDP ($w_n L_n$), along with parameter estimates for consumer expenditure shares (α_n^j), producer expenditure shares (γ_n^j, $\gamma_n^{k,j}$), and the trade elasticities (θ^j).

I obtain values for π_{ni}^j and $w_n L_n$ in the baseline year of 2014 from the 2016 version of the World Input Output Database (Timmer et al., 2015b). This database contains global input-output tables for 56 sectors and 43 countries and an aggregate region for the rest of the world. I aggregate this data into 13 countries/regions (including nine LMICs) and three sectors: agriculture and mining (referred to just as agriculture in the rest of the paper), manufacturing, and services. Appendix C.1 has further details on the aggregation scheme. I also use the World Input Output Database to calibrate the consumer and producer expenditure share parameters by setting them equal to the relevant expenditure shares observed in the data in 2014. Finally, I use estimates of θ^j from Caliendo and Parro (2015). Caliendo and Parro (2015) estimate θ^j for a number of sub-sectors within the sectors that I call agriculture and manufacturing. To obtain estimates of θ^j that match my sectoral aggregation, I take a weighted average of the disaggregated estimates using global GDP shares of the relevant sub-sectors as weights. They set a single trade elasticity of 8.22 for all service sectors, so I can map this directly to my model without having to do any aggregation.

I then implement the various counterfactuals as follows. Eliminating trade in sector k involves letting $\hat{d}_{ni}^k$ go to infinity for all $n \neq i$ and setting all other $\hat{d}_{ni}^j = 1$. Eliminating trade imbalances involves setting $D_n' = 0$ for all n. Finally, imposing complete autarky involves setting $D_n' = 0$ for all n, letting $\hat{d}_{ni}^j$ go to infinity for all $n \neq i$ and setting $\hat{d}_{ni}^j = 1$ otherwise.

All data used to calibrate the model is in current price terms and the output of the model is also in current price terms. Consequently, I discuss the impact of all counterfactual changes in terms of the impact on nominal value added shares by sector. Given the assumptions of a perfectly competitive labour market with a homogenous labour input and no adjustment frictions, wages are equalised across all sectors. The fact that labour is the only factor of production also implies that the labour share of value added is one in all sectors. It follows that nominal value added shares in the model are also equal to sectoral employment shares, which provides an addition reason to focus on the changes in nominal value added shares as the relevant output of the model.

3.4.3 Counterfactual Results

This section outlines the results from the various counterfactual exercises. I begin by eliminating trade in each of the three sectors of the model one at a time, while holding the ratio of each economy's trade deficit to global GDP constant. I then eliminate trade imbalances in the model holding trade costs in all sectors constant. These are the four basic counterfactuals. I then consider the effect of imposing these four counterfactual changes simultaneously (that is, imposing autarky on all economies). I also compute the value of the residual interaction effect from imposing the four basic counterfactuals simultaneously as the difference between the sum of the four changes from the individual counterfactuals and the effect of moving to autarky. In all cases, I focus on sectoral reallocations in nominal value added, with a particular focus on changes in the nominal value added share of manufacturing.

The aim of these exercises is to provide qualitative and quantitative insights on the effect of global integration on manufacturing shares in LMICs. Of course, the process of deepening and changing global integration since the 1960s described in Section 3.3 did not occur from an initial equilibrium of autarky, with trade in goods and services and trade imbalances all being features of the global economy even at that time. The counterfactual exercises do, however, provide a sense of what the upper bounds on the relevant effects might be. The advantage of imposing these counterfactuals is this that I do not have to estimate changes in trade costs since the 1960s to derive my results.

Columns (i)-(iii) of Table 3.4 show the changes in the nominal share of value added in manufacturing when eliminating all trade in agriculture, manufacturing and services, respectively[5]. As shown in Figures 3.10, 3.11 and 3.12, the baseline sectoral trade balances are closely correlated with the counterfactuals changes in manufacturing shares. In particular, trade surpluses in agriculture (services) are positively correlated with the change in manufacturing share when trade in agriculture (services) is eliminated. Conversely, trade surpluses in manufacturing are negatively correlated with the change in manufacturing share when trade in manufacturing is eliminated. These results are intuitive. When trade in a given sector is shut down, economies that run surpluses in that sector reallocate some of the resources from that sector to the other two sectors of the economy, while economies that run deficits in those sectors expand them at the expense of the other two sectors in the economy.

The slope of the regression line in Figures 3.10, 3.11 and 3.12 is less than one partly because trade balances are measured in terms of total expenditures/revenue, while the manufacturing share is measured in terms of value added. Because of the presence of intermediate inputs in the model, changes in the total value of trade flows are larger than the associated changes in value added from the reallocation of labour. The slope is also larger in absolute value when trade in manufactures is shut down, as the manufacturing sector is hit directly by the shock. In contrast, when trade in services is shut down, for example, countries that were running services sector surpluses shrink their service sector and split the reallocation of resources across both the agriculture and manufacturing sectors. Consequently, the regression slopes are approximately equal in magnitude for the agriculture and service sector counterfactuals, and roughly double the magnitude for the manufacturing sector counter-

[5]For a summary of the key features of the baseline equilibrium, see the relevant figures in Appendix C.4; full details of sectoral reallocations associated with each counterfactual can be found in Appendix C.5.

factual. These results show that while direct exposure to competition from manufacturing imports has been a cause of deindustrialisation pressures in many countries, the direct effect of trade in agriculture and services has potentially played a role than is of a similar order of magnitude.

Romania and Bulgaria provide examples of two small open economies that are have significant service export sectors. These economies ran substantial trade surpluses in services relative to GDP in 2014 (12% and 10.1%, respectively) and shutting down trade in services leads to 4.7 and 3.7 percentage point increases in their respective manufacturing shares. The evidence presented in Subsection 3.3.2 suggests that a substantial number of LMICs may run surpluses at least as large as these, a trend that may be further enhanced with future technological change.

Column (iv) of Table 3.4 shows the changes in the nominal share of value added in manufacturing when eliminating all trade imbalances in the global economy. As illustrated in Figure 3.13, the change in the manufacturing share is negatively correlated with the size of an economy's trade balance. Countries that run trade deficits (surpluses) tend to see an increase (decrease) in their manufacturing share and these effects are larger the bigger the trade imbalance. Although these results provide qualitative support for the notion that trade deficits in LMICs may be responsible for deindustrialisation pressures in these countries, the magnitude of these effects is small relative to the size of the other counterfactual effects. The two economies with the largest gain in manufacturing share were Bulgaria and the Rest of World region at 0.8 and 0.7 percentage points, respectively. Their deficit to GDP ratios in 2014 of -5.4% and -11.9%, respectively, are smaller than the average imbalance amongst deficit economies (see Subsection 3.3.3), so it is possible that eliminating imbalances could have a larger effect on the more extreme ends of the entire cross country distribution. That said, the evidence presented here is not supportive of the fact that these effects alone would be able to account for a large degree of deindustrialisation.

Column (vi) of Table 3.4 shows the changes in the nominal share of value added in manufacturing when implementing counterfactuals (i)-(iv) together (that is, imposing autarky). As illustrated in Figure 3.14, the change in the manufacturing share is negatively correlated with the size of an economy's trade balance in manufacturing. The only country that this qualitative pattern doesn't hold for is Turkey as it ran a surplus in all sectors in 2014; however, in this case, the sector with the smallest surplus is the one that shrunk. All LMICs other than China and India ran trade deficits in manufacturing in 2014 and all of these countries see an increase in their manufacturing share by moving to autarky. China and India are the only LMICs who lose manufacturing share by moving to autarky in this period.

Overall, the results of these counterfactuals provide qualitative support for the potential deindustrialisation mechanisms highlighted in Section 3.3. In the context of the model, openness to trade in manufactures has lead to manufacturing output across the LMICs being relatively more concentrated, while openness in agriculture and services has allowed economies without a comparative advantage in manufacturing to expand production in these sectors, further pulling resources away from manufacturing. Trade deficits have also added marginally to these effects. Quantitatively, however, none of these effects are large relative to the deindustrialisation trends observed in the data. For consistency, I run regression (3.1) on the country/region aggregation scheme used to generate the results in Table 3.4 and I find deindustrialisation trends of a similar magnitude to those in Section 3.2. In contrast, the average change in the manufacturing share across all LMICs in the sample from imposing autarky is just 1.1 percentage points (with none of the direct effects from the individual counterfactual changes exceeding 1.0 percentage points).

Table 3.4: Counterfactual Changes in Manufacturing Share of Nominal Value Added

	Baseline: VA Shares			Baseline: Trade Bal.			Change in Manufacturing Share					
	A	M	S	A	M	S	(i)	(ii)	(iii)	(iv)	(v)	**(vi)**
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	-2.4	2.4	2.3	0.8	-1.1	**1.9**
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	0.8	1.1	-0.6	0.3	-0.7	**0.9**
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	-1.0	-2.7	0.2	-0.7	2.0	**-2.2**
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	1.5	0.6	-0.5	0.0	-0.7	**1.0**
IND	16.7	16.6	66.7	-5.3	1.6	3.7	-1.5	-0.6	0.6	0.2	0.5	**-0.7**
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	1.7	0.9	-0.7	-0.1	-0.9	**0.9**
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	-1.1	3.3	4.3	0.3	-3.6	**3.3**
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	3.1	4.3	1.1	-0.3	-3.1	**5.1**
TUR	9.8	18.6	71.6	0.3	0.5	1.8	0.2	-0.2	0.6	-0.6	0.1	**0.1**
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	0.1	1.0	0.8	0.0	-0.8	**1.1**
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	-1.4	-2.2	0.3	-0.2	1.4	**-1.9**
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	-0.1	-1.3	1.3	-0.7	-0.1	**-0.9**
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	-0.1	1.9	0.8	0.6	-1.5	**1.8**
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	1.4	1.9	-2.5	0.7	-0.2	**1.2**
WLD	8.9	17.0	74.1	0.0	0.0	0.0	0.0	-0.2	0.2	0.0	0.0	**0.1**

Notes: The first set of columns show baseline nominal value added shares by sector (as percentages); the second set of columns show the baseline trade balance by sector (as a percentage of GDP). The final set of columns shows the percentage point change in the manufacturing share of value added associated with the following counterfactuals: (i) no trade in agriculture; (ii) no trade in manufacturing; (iii) no trade in services; (iv) no trade imbalances; (v) interaction effects from (i)-(iv); (vi) autarky. The row titled LMIC Ave. presents the the average values across all LMICs for each column.

Figure 3.10: Change in Nominal Manufacturing Share from Eliminating Trade in Agriculture

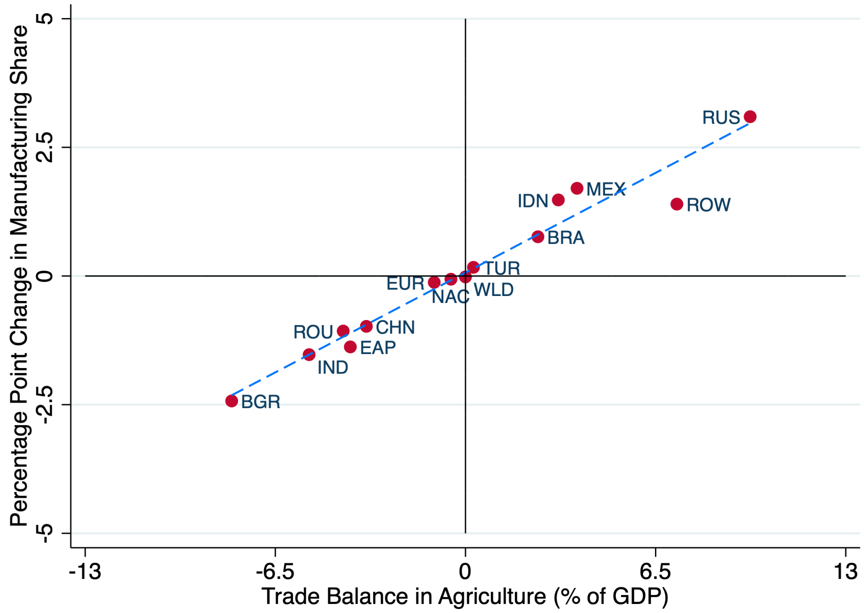

Notes: This figure plots each country/region's trade balance in agriculture as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the elimination of trade in agriculture.

Figure 3.11: Change in Nominal Manufacturing Share from Eliminating Trade in Manufacturing

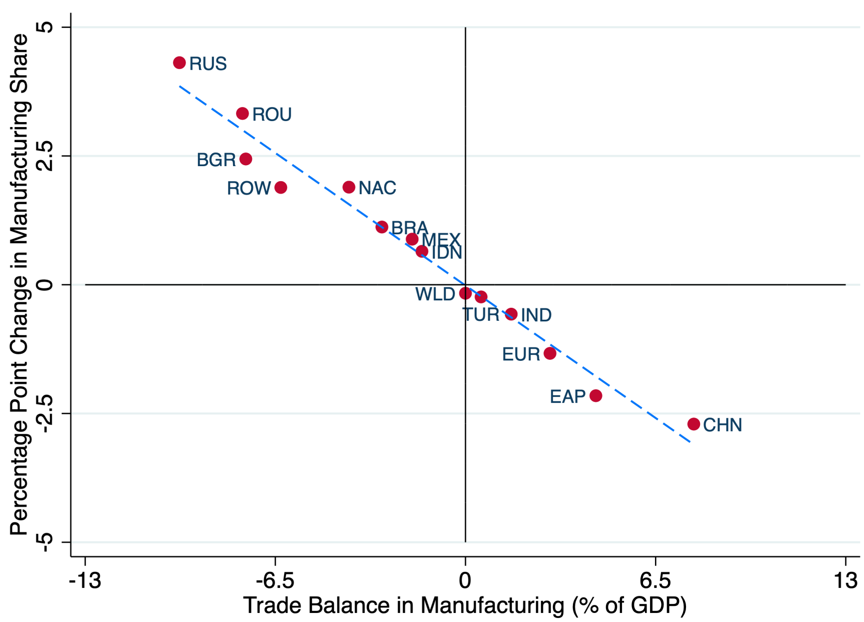

Notes: This figure plots each country/region's trade balance in manufacturing as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the elimination of trade in manufacturing.

Figure 3.12: Change in Nominal Manufacturing Share from Eliminating Trade in Services

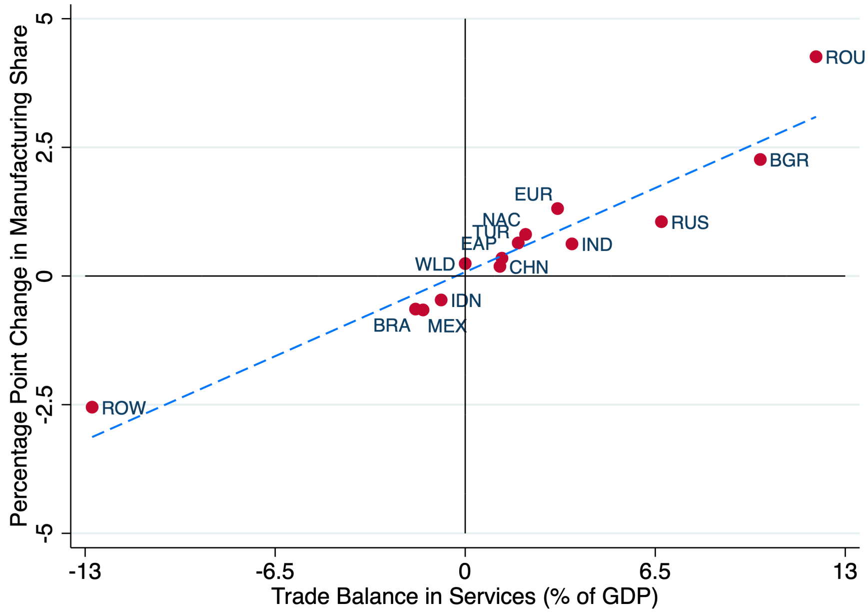

Notes: This figure plots each country/region's trade balance in services as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the elimination of trade in services.

Figure 3.13: Change in Nominal Manufacturing Share from Eliminating Trade Imbalances

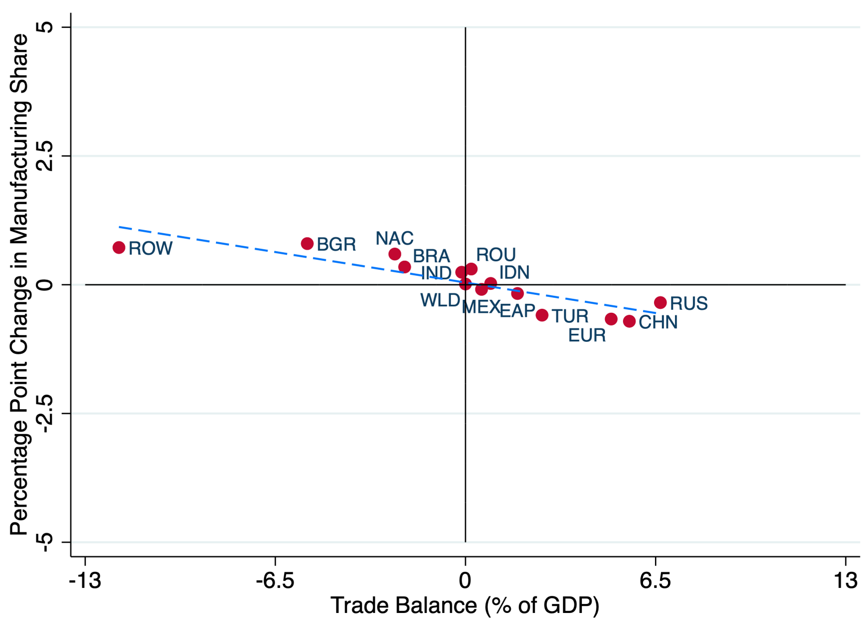

Notes: This figure plots each country/region's trade balance as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the elimination of trade imbalances.

Figure 3.14: Change in Nominal Manufacturing Share from Interaction Effects

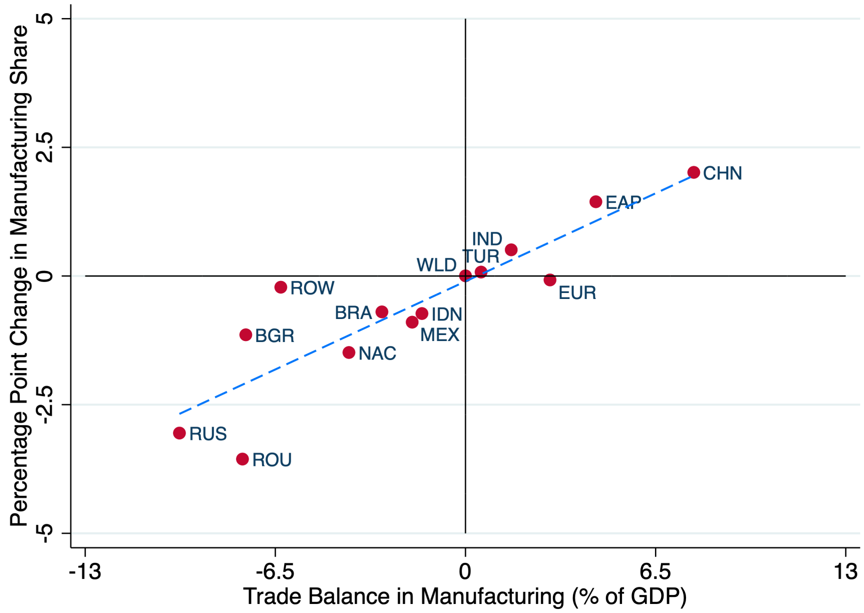

Notes: This figure plots each country/region's trade balance in manufacturing as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the interaction effects from shutting down each dimension of trade.

Figure 3.15: Change in Nominal Manufacturing Share from Eliminating all Trade

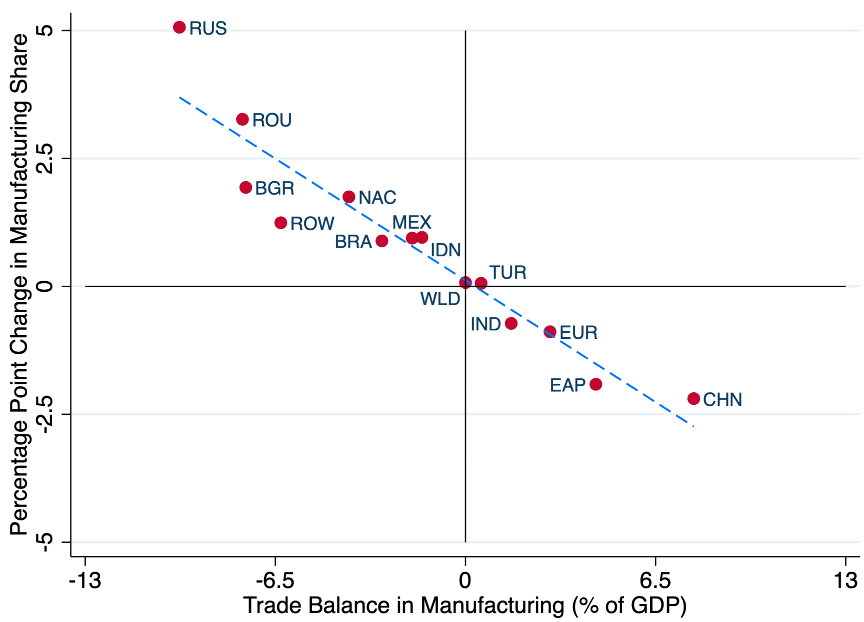

Notes: This figure plots each country/region's trade balance in manufacturing as a percentage of their GDP on the horizontal axis against the percentage point change in the manufacturing share of nominal value added on the vertical axis caused by the elimination of all trade.

3.5 Conclusion

The results presented in this chapter suggest that changing patterns of global integration may be partly responsible for the observed patterns of premature deindustrialisation. While closing down trade in the model has almost no effect on global manufacturing shares, it does lead to reallocations of manufacturing output that increase the average manufacturing share across the nine LMICs I study. Although the counterfactual changes are qualitatively in line with the potential channels of deindustrialisation proposed in Section 3.3, the magnitude of these changes tends to be small with an average gain in manufacturing share of just 1.1 percentage points across the nine LMICs from moving to autarky.

It is worth noting at least two caveats to these results. First, the model presented in this chapter fails to admit two key closed economy channels of structural transformation: the Baumol and Engel effects. There is reason to think that moving away from Cobb-Douglas utility to a utility function with a lower elasticity of substitution in consumption across sectors would amplify the changes in manufacturing shares in the counterfactual exercises. A lower elasticity of substitution across sectors would mean that consumers are less willing to absorb trade imbalances through changes in domestic consumption patterns when trade is shut down, leading to larger reallocations across sectors. Indeed, the fact that China's nominal manufacturing share in autarky is predicted to be just 2.2 percentage points lower than the baseline value of 30.1% in 2014 (see Table 3.4) is an example of a quantitatively surprising implication of the model that could be resolved by such a change. For countries that are deindustrialising due to income effects, adding non-homotheticicities to the model would lead to higher manufacturing shares when trade is shut down given the positive effect of trade on incomes. Adding this effect could, therefore, also increase the magnitude of the changes in manufacturing shares. It would be useful to extend the model to accommodate these mechanisms so that these hypotheses can be formally tested.

A second caveat is that the LMICs studied in this chapter are not representative of all LMICs given that the WIOD typically contains data for larger, more prosperous economies. The three economies that saw the biggest gain in manufacturing share on average across the counterfactual exercises were Bulgaria, Romania and Russia. These economies, particularly Bulgaria and Romania, have been particularly exposed to a number of the globalisation trends identified in Section 3.3: they are small, run large trade deficits relative to GDP, and run large trade surpluses in services and large deficits in manufacturing. This suggests that the magnitude of the effects may have been larger with a more representative sample of economies. A general neglect of low- and lower-middle-income countries due to data constraints is a major shortcoming across the quantitative structural transformation literature.

Assessing the industrialisation experiences of these countries at the more extreme end of the cross country income distribution through the lens of our models is likely to be an important part of making sense of the data presented in Section 3.2.

Finally, it is worth noting that in addition to being premature in a positive sense, Rodrik (2016) also refers to LMIC deindustrialisation trends as being premature in a normative sense given the view that the manufacturing sector has a special role to play in the development process (Rodrik, 2008, 2013). While this may be the case, the theoretical framework I use ascribes no special role to manufacturing. Moreover, there are no market imperfections in the model and the decentralised equilibrium corresponds to the planner's solution at the global level. It thus follows that while the framework I use in this paper provides some insight into the drivers of the industrialisation trends observed in LMICs, it provides few insights into debates about the potential consequences and appropriate policy responses. Given the concern that many academics and policymakers in LMICs have regarding premature deindustrialisation, it would be useful to develop more models with the kinds of imperfections that accommodate the possibility that industrialisation is a necessary condition for economic development. Taking these models to the data would give more policy-relevant insights into the consequences of the industrialisation trends discussed in this paper and how policymakers ought to respond.

3.6 References

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of political Economy*, 116(3):467–498.

Acosta, P. A., Lartey, E. K., and Mandelman, F. S. (2009). Remittances and the dutch disease. *Journal of international economics*, 79(1):102–116.

Alessandria, G. A., Johnson, R. C., and Yi, K.-M. (2021). Perspectives on trade and structural transformation. *National Bureau of Economic Research Working Paper*.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Buera, F. J. and Kaboski, J. P. (2009). Can traditional theories of structural change fit the data? *Journal of the European Economic Association*, 7(2-3):469–477.

Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies*, 82(1):1–44.

Corden, W. M. and Neary, J. P. (1982). Booming sector and de-industrialisation in a small open economy. *The economic journal*, 92(368):825–848.

Cravino, J. and Sotelo, S. (2019). Trade-induced structural change and the skill premium. *American Economic Journal: Macroeconomics*, 11(3):289–326.

Dasgupta, S. and Singh, A. (2007). Manufacturing, services and premature deindustrialization in developing countries: A kaldorian analysis. In *Advancing development*, pages 435–454. Springer.

Dekle, R., Eaton, J., and Kortum, S. (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540.

Duarte, M. and Restuccia, D. (2010). The role of the structural transformation in aggregate productivity. *The Quarterly Journal of Economics*, 125(1):129–173.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the penn world table. *American economic review*, 105(10):3150–82.

Felipe, J., Mehta, A., and Rhee, C. (2019). Manufacturing matters...but its the jobs that count. *Cambridge Journal of Economics*, 43(1):139–168.

Foellmi, R. and Zweimüller, J. (2008). Structural change, engel's consumption cycles and kaldor's facts of economic growth. *Journal of monetary Economics*, 55(7):1317–1328.

Fujiwara, I. and Matsuyama, K. (2020). A technology-gap model of premature deindustrialization. *Available at SSRN 3753930*.

Gollin, D., Jedwab, R., and Vollrath, D. (2016). Urbanization with and without industrialization. *Journal of Economic Growth*, 21(1):35–70.

Haraguchi, N., Cheng, C. F. C., and Smeets, E. (2017). The importance of manufacturing in economic development: has this changed? *World Development*, 93:293–315.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2014). Growth and structural transformation. *Handbook of economic growth*, 2:855–941.

Huneeus, F. and Rogerson, R. (2020). Heterogeneous paths of industrialization. Technical report, National Bureau of Economic Research.

Klasing, M. J. and Milionis, P. (2014). Quantifying the evolution of world trade, 1870–1949. *Journal of International Economics*, 92(1):185–197.

Kruse, H., Mensah, E., Sen, K., and de Vries, G. (2021). A manufacturing renaissance? industrialization trends in the developing world. Technical report, WIDER Working Paper.

Lartey, E. K., Mandelman, F. S., and Acosta, P. A. (2012). Remittances, exchange rate regimes and the dutch disease: A panel data analysis. *Review of International Economics*, 20(2):377–395.

Lewis, L. T., Monarch, R., Sposi, M., and Zhang, J. (2022). Structural change and global trade. *Journal of the European Economic Association*, 20(1):476–512.

Matsuyama, K. (2009). Structural change in an interdependent world: A global view of manufacturing decline. *Journal of the European Economic Association*, 7(2-3):478–486.

Matsuyama, K. (2019). Engel's law in the global economy: Demand-induced patterns of structural change, innovation, and trade. *Econometrica*, 87(2):497–528.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American economic review*, 97(1):429–443.

Nguimkeu, P. and Zeufack, A. (2019). Manufacturing in structural change in africa. *World Bank Policy Research Working Paper*.

Our World in Data (2022). Trade and Globalization. *https://ourworldindata.org/trade-and-globalization Accessed on 12 July, 2022*.

Palma, G. et al. (2005). Four sources of de-industrialisation and a new concept of the dutch disease. *Beyond reforms: structural dynamics and macroeconomic vulnerability*, 3(5):71–116.

Palma, J. G. (2014). De-industrialisation, premature de-industrialisation and the dutch-disease. *Revista NECAT-Revista do Núcleo de Estudos de Economia Catarinense*, 3(5):7–23.

Rajan, R. G. and Subramanian, A. (2011). Aid, dutch disease, and manufacturing growth. *Journal of development Economics*, 94(1):106–118.

Rodrik, D. (2008). The real exchange rate and economic growth. *Brookings papers on economic activity*, 2008(2):365–412.

Rodrik, D. (2013). Unconditional convergence in manufacturing. *The quarterly journal of economics*, 128(1):165–204.

Rodrik, D. (2016). Premature deindustrialization. *Journal of economic growth*, 21(1):1–33.

Sposi, M. (2019). Evolving comparative advantage, sectoral linkages, and structural change. *Journal of Monetary Economics*, 103:75–87.

Sposi, M., Yi, K.-M., and Zhang, J. (2021). Deindustrialization and industry polarization. Technical report, National Bureau of Economic Research.

Święcki, T. (2017). Determinants of structural change. *Review of Economic Dynamics*, 24:95–131.

Teignier, M. (2018). The role of trade in structural transformation. *Journal of Development Economics*, 130:45–65.

Timmer, M., de Vries, G. J., and De Vries, K. (2015a). Patterns of structural change in developing countries. In *Routledge handbook of industry and development*, pages 79–97. Routledge.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and De Vries, G. J. (2015b). An illustrated user guide to the world input–output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.

UNCTAD STAT (2022). Data Centre. *https://unctadstat.unctad.org/wds/ Accessed on 13 July, 2022.*

United Nations (2021). National accounts - analysis of main aggregates (ama). Technical report.

Uy, T., Yi, K.-M., and Zhang, J. (2013). Structural change in an open economy. *Journal of Monetary Economics*, 60(6):667–682.

Wise, A. (2021). *Structural Change in a Global Economy*. PhD thesis, Princeton University.

World Bank (2022). DataBank. *https://databank.worldbank.org/home.aspx Accessed on 13 July, 2022.*

Appendix A

Appendices to Chapter 1

A.1 Laissez-Faire Labour Allocation

The composite energy producer's first order conditions for dirty and green energy demand are given by

$$p_{e,t}E_{y,t}^{\frac{1}{\sigma_e}}\omega_k E_{k,t}^{\frac{-1}{\sigma_e}} = p_{k,t} \quad k \in \{d, g\}$$

For the dirty and green energy producers to demand a positive but finite amount of labour, it must be the case that

$$p_{k,t} = \frac{w_t}{A_{k,t}} \quad k \in \{d, g\}$$

Combining the two first order conditions for the composite energy producer and the pricing functions for dirty and green energy implies

$$E_{k,t} = \left(\frac{\omega_k A_{k,t}}{\omega_j A_{j,t}}\right)^{\sigma_e} E_{j,t}$$

for $j, k \in \{d, g\}$. Plugging this expression into the production function for composite energy give us

$$\begin{aligned} E_{y,t} &= \Gamma_{k,t}E_{k,t} \quad k \in \{d, g\} \\ \Gamma_{k,t} &\equiv \left[\omega_d^{\sigma_e}A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e}A_{g,t}^{\sigma_e-1}\right]^{\frac{\sigma_e}{\sigma_e-1}} (\omega_k A_{k,t})^{-\sigma_e} \quad j \in \{d, g\} \end{aligned}$$

The first order conditions for the final goods producer's choice of labour and energy inputs are

$$\begin{aligned} w_t &= (1-\nu-\alpha)\frac{Y_t}{L_{y,t}} \\ p_{e,t} &= \nu\frac{Y_t}{E_{y,t}} \end{aligned}$$

Combining these with the conditions above, we have

$$(1-\nu-\alpha)\frac{Y_t}{L_{y,t}} = \nu\frac{Y_t}{E_{y,t}}E_{y,t}^{\frac{1}{\sigma_e}}E_{k,t}^{\frac{-1}{\sigma_e}}A_{k,t} \quad k \in \{d, g\}$$

Plugging in the expression for $E_{y,t}$ and simplifying:

$$L_{k,t} = \left(\frac{\nu}{1-\nu-\alpha}\right)\left(\frac{\omega_k^{\sigma_e} A_{k,t}^{\sigma_e-1}}{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e-1}}\right) L_{y,t} \quad k \in \{d, g\}$$

Plugging this into the labour market clearing condition then gives us

$$\begin{aligned} L_{y,t} &= \frac{1-\nu-\alpha}{1-\alpha} N_t^{t-1} \\ \Rightarrow L_{y,t} &= \frac{\nu}{1-\alpha}\left(\frac{\omega_k^{\sigma_e} A_{k,t}^{\sigma_e-1}}{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e-1}}\right) N_t^{t-1} \end{aligned}$$

A.2 Lemma 1: Setup of Optimal Carbon Tax Problem

The social welfare function can be rewritten as

$$\mathbb{E}_t\left[U_t\right] = \sum_{s=t}^{\infty} \phi^{s-t}\mathbb{E}_t\left[N_s^s \ln c_s^s + \left(\frac{\beta}{\phi}\right) N_s^{s-1} \ln c_s^{s-1} + \left(\frac{\beta}{\phi}\right)^2 N_s^{s-2} \ln c_s^{s-2}\right]$$

By assumption, $c_s^s = \theta_s^y c_s^{s-1}$. Furthermore, given that the ratio of retiree to middle age income is $\frac{\alpha}{1-\alpha}$ and middle-aged households consume a fraction $\frac{1}{1+\psi_{s+1}^{s-1}\beta}$ of their earnings, we have

$$c_s^{s-2} = \frac{\alpha\left(1+\psi_{s+1}^{s-1}\beta\right)\psi_s^{s-1}\left(1+n_{s-1}\right)}{\left(1-\alpha\right)\psi_s^{s-2}\psi_{s-1}^{s-2}} c_{h,s}$$

from which it follows that

$$\begin{aligned} c_s^{s-2} &= \theta_s^o c_s^{s-1} \\ \theta_s^o &\equiv \frac{\alpha\left(1+\psi_{s+1}^{s-1}\beta\right)\left(1+n_{s-1}\right)\left[\psi_s^{s-1}+\theta_s^y\left(1+n_s\right)\right]}{\left(1-\alpha\right)\psi_s^{s-2}\psi_{s-1}^{s-2}} \end{aligned}$$

The social welfare function is thus

$$\begin{aligned} \mathbb{E}_t\left[U_t\right] &= \sum_{s=t}^{\infty} \phi^{s-t}\mathbb{E}_t\left[\tilde{N}_s \ln c_s^{s-1}\right] + \mathbb{E}_t\left[\Omega_t\right] \\ \tilde{N}_s &\equiv N_s^s + \left(\frac{\beta}{\phi}\right) N_s^{s-1} + \left(\frac{\beta}{\phi}\right)^2 N_s^{s-2} \\ \Omega_t &\equiv \sum_{s=t}^{\infty} \phi^{s-t}\left[N_s^s \ln \theta_s^y + \left(\frac{\beta}{\phi}\right)^2 N_s^{s-2} \ln \theta_s^o\right] \end{aligned}$$

We can then write the primal form of the policymaker's problem by eliminating prices from the relevant constraints and writing them purely in terms of the allocations. In particular, in addition to the technological and market clearing constraints of the economy, the policymaker is bound by the following competitive equilibrium conditions

$$\begin{aligned} c_s^s &= \theta_s^y c_s^{s-1} \\ c_s^{s-2} &= \theta_s^o c_s^{s-1} \\ k_{s+1}^{s-1} N_s^{s-1} &= s_s Y_s \\ s_s &= \frac{\psi_{s+1}^{s-1}\beta\left(1-\alpha\right)}{1+\psi_{s+1}^{s-1}\beta} \end{aligned}$$

These conditions combined with the goods market clearing condition imply

$$\begin{aligned}(1-s_s)\,Y_s &= c_s^{s-1}\bar{N}_s \\ \bar{N}_s &\equiv \theta_s^y N_s^s + N_s^{s-1} + \theta_s^o N_s^{s-2}\end{aligned}$$

The optimal carbon tax problem in primal form is, therefore, for the policymaker to choose a sequence of allocations $\left\{c_s^{s-1}, k_{s+1}^{s-1}, E_{y,s}, K_{y,s}, L_{y,s}, E_{d,s}, E_{g,s}, L_{d,s}, L_{g,s}, S_s\right\}$ to maximise

$$\sum_{s=t}^{\infty} \phi^{s-t}\mathbb{E}_t\left[\tilde{N}_s \ln c_s^{s-1}\right]$$

subject to

$$\begin{aligned}c_s^{s-1}\bar{N}_s + K_{y,s+1} &= e^{-\gamma_s\left(S_s-\bar{S}\right)}A_{y,s}E_{y,s}^{\nu}K_{y,s}^{\alpha}L_{y,s}^{1-\nu-\alpha} \\ c_s^{s-1}\bar{N}_s &= (1-s_s)\,e^{-\gamma_s\left(S_s-\bar{S}\right)}A_{y,s}E_{y,s}^{\nu}K_{y,s}^{\alpha}L_{y,s}^{1-\nu-\alpha} \\ E_{y,s} &= \left(\omega_d E_{d,s}^{\frac{\sigma_e-1}{\sigma_e}} + \omega_g E_{g,s}^{\frac{\sigma_e-1}{\sigma_e}}\right)^{\frac{\sigma_e-1}{\sigma_e}} \\ E_{k,s} &= A_{k,s}L_{k,s} \quad k\in\{d,g\} \\ N_s^{s-1} &= L_{y,s} + L_{d,s} + L_{g,s} \\ S_s - \bar{S} &= \sum_{i=0}^{s}(1-d_i)\,E_{d,s-i}\end{aligned}$$

Note that since Ω_t, which measures the impact of intratemporal inequality between generations on social welfare, is not a function of any of the choice variables of the problem, this term can be dropped from the objective function.

A.3 Proposition 1: Derivation of Optimal Carbon Tax

The Lagrangian is then given by

$$\sum_{s=t}^{\infty}\phi^{s-t}\mathbb{E}_t\left[\tilde{N}_s\ln c_s^{s-1}-\lambda_{y,s}\left\{c_s^{s-1}\bar{N}_s+K_{y,s+1}-e^{-\gamma_s\left(S_s-\bar{S}\right)}A_{y,s}E_{y,s}^{\nu}K_{y,s}^{\alpha}L_{y,s}^{1-\nu-\alpha}\right\}\right.$$
$$-\mu_s\left\{c_s^{s-1}\bar{N}_s-(1-s_s)\,e^{-\gamma_s\left(S_s-\bar{S}\right)}A_{y,s}E_{y,s}^{\nu}K_{y,s}^{\alpha}L_{y,s}^{1-\nu-\alpha}\right\}$$
$$-\lambda_{e,s}\left\{E_{y,s}-\left(\omega_d E_{d,s}^{\frac{\sigma_e-1}{\sigma_e}}+\omega_g E_{g,s}^{\frac{\sigma_e-1}{\sigma_e}}\right)^{\frac{\sigma_e-1}{\sigma_e}}\right\}-\sum_{k\in\{d,g\}}\lambda_{k,s}\left\{E_{k,s}-A_{k,s}L_{k,s}\right\}$$
$$\left.-\lambda_{l,s}\left\{L_{y,s}+L_{d,s}+L_{g,s}-N_s^{s-1}\right\}-\varepsilon_s\left\{\sum_{i=0}^{s}(1-d_i)\,E_{d,s-i}-\left(S_s-\bar{S}\right)\right\}\right]$$

The first order condition on c_s^{s-1} is

$$\tilde{N}_s\frac{1}{c_s^{s-1}}=(\lambda_{y,s}+\mu_s)\,\bar{N}_s$$
$$\Rightarrow\mu_s=\frac{\tilde{N}_s}{c_s^{s-1}\bar{N}_s}-\lambda_{y,s}$$
$$\Rightarrow\lambda_{y,s+1}=\frac{\tilde{N}_{s+1}}{c_{s+1}^{s}\bar{N}_{s+1}}-\mu_{s+1}$$

The first oder condition on $K_{y,s+1}$ is

$$\lambda_{y,s}=\mathbb{E}_s\left[\phi\alpha\frac{Y_{s+1}}{K_{y,s+1}}\left\{\lambda_{y,s+1}+\mu_{s+1}(1-s_{s+1})\right\}\right]$$

Plugging the expression for $\lambda_{y,s}$ into the one for μ_s and then plugging the expression for $\lambda_{y,s+1}$ into the resulting equation gives us

$$\mu_s=\frac{\tilde{N}_s}{c_s^{s-1}\bar{N}_s}-\mathbb{E}_s\left[\phi\alpha\frac{Y_{s+1}}{K_{y,s+1}}\left\{\lambda_{y,s+1}+\mu_{s+1}(1-s_{s+1})\right\}\right]$$
$$\Rightarrow\mu_s=\mathbb{E}_s\left[\Lambda_{s+1}\right]+\mathbb{E}_s\left[\eta_{s+1}\mu_{s+1}\right]$$
$$\Lambda_{s+1}\equiv\frac{\tilde{N}_s}{c_s^{s-1}\bar{N}_s}-\phi\alpha\frac{Y_{s+1}}{K_{y,s+1}}\frac{\tilde{N}_{s+1}}{c_{s+1}^{s}\bar{N}_{s+1}}$$
$$\eta_{s+1}\equiv\phi\alpha\frac{Y_{s+1}}{K_{y,s+1}}s_{s+1}$$

Assuming that

$$\lim_{i\to\infty}\mathbb{E}_s\left[\left(\prod_{q=s}^{i-1}\eta_{q+1}\right)\Lambda_{i+1}\right]=0$$

the difference equation for μ_s then implies

$$\mu_s = \sum_{i=s}^{\infty} \mathbb{E}_s \left[\left(\prod_{q=s}^{i-1} \eta_{q+1} \right) \Lambda_{i+1} \right]$$

Plugging the expression for η_{q+1} into the $\left(\prod_{q=s}^{i-1} \eta_{q+1} \right)$ term from the μ_s equation above implies

$$\prod_{q=s}^{i-1} \eta_{q+1} = (\alpha\phi)^{i-s} \frac{s_i Y_i}{s_s Y_s}$$

Simplifying the equation for Λ_{i+1}

$$\Lambda_{i+1} = \frac{1}{Y_i} \left[\frac{1}{1-s_i} \tilde{N}_i - \frac{\phi\alpha}{s_i (1 - s_{i+1})} \tilde{N}_{i+1} \right]$$

Plugging these expressions for $\prod_{q=s}^{i-1} \eta_{q+1}$ and Λ_{i+1} into the equation for μ_s gives us

$$\begin{aligned}
\mu_s =& \sum_{i=s}^{\infty} \mathbb{E}_s \left[\left\{ (\alpha\phi)^{i-s} \frac{s_i Y_i}{s_s Y_s} \right\} \frac{1}{Y_i} \left\{ \frac{1}{1-s_i} \tilde{N}_i - \frac{\phi\alpha}{s_i (1 - s_{i+1})} \tilde{N}_{i+1} \right\} \right] \\
\Rightarrow \mu_s =& \frac{\mathbb{E}_s [\Theta_s]}{s_s Y_s} \\
\Theta_s \equiv& \sum_{i=s}^{\infty} (\alpha\phi)^{i-s} \left(\frac{s_i}{1-s_i} \tilde{N}_i - \frac{\phi\alpha}{1-s_{i+1}} \tilde{N}_{i+1} \right)
\end{aligned}$$

μ_s represents the shadow value for the planner of relaxing the savings constraint imposed in competitive equilibrium that requires $K_{y,s+1} = s_s Y_s$. When the planner is in a first best scenario where they can optimise both the carbon tax and the aggregate savings rate, we would have $\mu_s = 0$ (since the savings constraint would not be binding), which happens when $\mathbb{E}_s [\Lambda_{s+1}] = 0$. This result should be intuitive since $\mathbb{E}_s [\Lambda_{s+1}] = 0$ happens when

$$\frac{\tilde{N}_s}{c_s^{s-1} \bar{N}_s} = \mathbb{E}_s \left[\phi\alpha \frac{Y_{s+1}}{K_{y,s+1}} \frac{\tilde{N}_{s+1}}{c_{s+1}^{s} \bar{N}_{s+1}} \right]$$

which would be the Euler equation on the choice of capital chosen by the planner in the first best scenario.

Now consider the first order conditions on S_s and $E_{d,s}$:

$$\varepsilon_s = \lambda_{y,s}\gamma_s Y_s + \mu_s \left(1 - s_s\right) \gamma Y_s$$

$$\lambda_{e,s} E_{y,s}^{\frac{1}{\sigma_e}} \omega_d E_{d,s}^{\frac{-1}{\sigma_e}} = \lambda_{d,s} + \sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\epsilon_{s+i} \left(1 - d_i\right)\right]$$

Starting by simplifying the expression

$$\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\epsilon_{s+i} \left(1 - d_i\right)\right]$$

Combining this with the expression for ϵ_s gives us

$$\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\epsilon_{s+i} \left(1 - d_i\right)\right] = \sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\left(1 - d_i\right) \gamma_{s+i} Y_{s+i} \left\{\lambda_{y,s+i} + \mu_{s+i} \left(1 - s_{s+i}\right)\right\}\right]$$

From the first oder conditions on c_s^{s-1}, we can write this as

$$\begin{aligned}
&\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\left(1 - d_i\right) \gamma Y_{s+i} \left\{\lambda_{y,s+i} + \mu_{s+i} \left(1 - s_{s+i}\right)\right\}\right] \\
=&\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\left(1 - d_i\right) \gamma_{s+i} Y_{s+i} \left\{\lambda_{y,s+i} + \mu_{s+i} - \mu_{s+i} s_{s+i}\right\}\right] \\
=&\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\left(1 - d_i\right) \gamma_{s+i} Y_{s+i} \left(\frac{\tilde{N}_{s+i}}{c_{s+i}^{s} \bar{N}_{s+i}} - \mu_{s+i} s_{s+i}\right)\right] \\
=&\sum_{i=0}^{\infty} \phi^i \mathbb{E}_s \left[\left(1 - d_i\right) \gamma_{s+i} \left\{\frac{\tilde{N}_{s+i}}{\left(1 - s_{s+i}\right)} - \Theta_{s+i}\right\}\right]
\end{aligned}$$

From FOC on $E_{y,s}$, $L_{y,s}$ and $L_{d,s}$

$$\begin{aligned}
\lambda_{e,s} &= \left[\lambda_{y,s} + \mu_s \left(1 - s_s\right)\right] \nu \frac{Y_s}{E_{y,s}} \\
\lambda_{l,s} &= \left[\lambda_{y,s} + \mu_s \left(1 - s_s\right)\right] \left(1 - \nu - \alpha\right) \frac{Y_s}{L_{y,s}} \\
\lambda_{l,s} &= \lambda_{d,s} A_{d,s}
\end{aligned}$$

Putting this together

$$\begin{aligned}
\left[\lambda_{y,s}+\mu_s\left(1-s_s\right)\right]\nu\frac{Y_s}{E_{y,s}}E_{y,s}^{\frac{1}{\sigma_e}}\omega_d E_{d,s}^{\frac{-1}{\sigma_e}} &= \left[\lambda_{y,s}+\mu_s\left(1-s_s\right)\right]\frac{\left(1-\nu-\alpha\right)}{A_{d,s}}\frac{Y_s}{L_{y,s}} \\
&\quad+\sum_{i=0}^{\infty}\phi^i\mathbb{E}_s\left[\left(1-d_i\right)\gamma_{s+i}\left\{\frac{\tilde{N}_{s+i}}{\left(1-s_{s+i}\right)}-\Theta_{s+i}\right\}\right] \\
\nu\frac{Y_s}{E_{y,s}}E_{y,s}^{\frac{1}{\sigma_e}}\omega_d E_{d,s}^{\frac{-1}{\sigma_e}} &= \frac{\left(1-\nu-\alpha\right)}{A_{d,s}}\frac{Y_s}{L_{y,s}} \\
&\quad+Y_s\sum_{i=0}^{\infty}\phi^i\mathbb{E}_s\left[\frac{\tilde{N}_{s+i}}{\tilde{N}_s}\left(1-d_i\right)\gamma_{s+i}\frac{\left(1-s_s\right)}{\left(1-s_{s+i}\right)}\frac{\left\{1-\frac{\left(1-s_{s+i}\right)\Theta_{s+i}}{\tilde{N}_{s+i}}\right\}}{\left\{1-\frac{\left(1-s_s\right)\mathbb{E}_s\left[\Theta_s\right]}{\tilde{N}_s}\right\}}\right]
\end{aligned}$$

This implies that the second best carbon tax to GDP ratio is given by

$$\frac{\tau_t^*}{Y_t^*}=\sum_{s=t}^{\infty}\phi^{s-t}\mathbb{E}_t\left[\frac{\tilde{N}_s}{\tilde{N}_t}\left(1-d_{s-t}\right)\gamma_s\frac{\left(1-s_t\right)}{\left(1-s_s\right)}\frac{\left\{1-\frac{\left(1-s_s\right)\Theta_s}{\tilde{N}_s}\right\}}{\left\{1-\frac{\left(1-s_t\right)\mathbb{E}_t\left[\Theta_t\right]}{\tilde{N}_t}\right\}}\right]$$

A.4 R_t and Demographic Change

This appendix illustrates the sufficient conditions for demographic change (defined as rising ψ_t^{t-1} and ψ_t^{t-2} along with falling n_{t-1}) to lead to a fall in R_t in the laissez-faire equilibrium. From the firm's first order condition on capital choice,

$$R_t = \alpha \frac{Y_t}{K_{y,t}}$$

The consumer and firm conditions imply that the capital stock is given by

$$K_{y,t} = (1-\alpha)\left(\frac{\psi_t^{t-2}\beta}{1+\psi_t^{t-2}\beta}\right) Y_{t-1}$$

The labour allocation to the final goods sector is

$$L_{y,t} = \left(\frac{1-\nu-\alpha}{1-\alpha}\right) N_t^{t-1}$$

Using the labour allocations to the clean and dirty energy varieties, the composite energy good can be written as

$$E_{y,t} = \left(\frac{\nu}{1-\alpha}\right) N_t^{t-1}$$

while the atmospheric carbon concentration can be written as

$$\begin{aligned}
S_t &= \Upsilon_t + \Delta_t N_t^{t-1} \\
\Upsilon_t &\equiv S_{1,t-1} + (1-\varphi) S_{2,t-1} \\
\Delta_t &\equiv [\varphi_L + (1-\varphi_L)\varphi_0] \left(\frac{\nu}{1-\alpha}\right) \left(\frac{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e - 1}}{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e-1}}\right)
\end{aligned}$$

where $S_{1,t-1}$ is the permanent component of S_{t-1} and $S_{2,t-1}$ is the transitory component of S_{t-1}. The real rate can then be written as

$$R_t = \tilde{A}_{y,t} e^{-\gamma\left(\Upsilon_t + \Delta_t N_t^{t-1} - \bar{S}\right)} \left[\left(\frac{1+\psi_t^{t-2}\beta}{\psi_t^{t-2}\beta}\right) \frac{N_t^{t-1}}{Y_{t-1}}\right]^{1-\alpha}$$

where

$$\tilde{A}_{y,t} \equiv \alpha A_{y,t} \left(\frac{\nu}{1-\alpha}\right)^{\nu} \left(\frac{1-\nu-\alpha}{1-\alpha}\right)^{1-\nu-\alpha} (1-\alpha)^{1-\alpha}$$

It follows that

$$\frac{dR_t}{d\psi_t^{t-2}} < 0$$

and

$$\frac{dR_t}{dN_t^{t-1}} > 0 \;\text{ iff }\; \frac{1-\alpha}{N_t^{t-1}} > \gamma\left[\varphi_L + (1-\varphi_L)\,\varphi_0\right]\left(\frac{\nu}{1-\alpha}\right)\left(\frac{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1}}{\omega_d^{\sigma_e} A_{d,t}^{\sigma_e-1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e-1}}\right)$$

Furthermore, since $N_t^{t-1} = \psi_t^{t-1}\,(1+n_{t-1})\,N_{t-2}^{t-2}$,

$$dN_t^{t-1} < 0 \;\text{ iff }\; \frac{d\psi_t^{t-1}}{\psi_t^{t-1}} < -\frac{dn_{t-1}}{1+n_{t-1}}$$

Appendix B

Appendices to Chapter 2

B.1 Country and Sector Names

Table B.1: Country Names

Country Index	Country Code	Country Name
1	AUS	Australia
2	AUT	Austria
3	BEL	Belgium
4	BGR	Bulgaria
5	CAN	Canada
6	CHE	Switzerland
7	CYP	Cyprus
8	CZE	Czech Republic
9	DNK	Denmark
10	ESP	Spain
11	EST	Estonia
12	FIN	Finland
13	FRA	France
14	GBR	United Kingdom
15	GRC	Greece
16	HRV	Croatia
17	HUN	Hungary
18	IRL	Ireland
19	ITA	Italy
20	KOR	Korea, Rep.
21	LTU	Lithuania
22	LUX	Luxembourg
23	LVA	Latvia
24	MEX	Mexico
25	MLT	Malta
26	NLD	Netherlands
27	NOR	Norway
28	POL	Poland
29	PRT	Portugal
30	ROU	Romania
31	SVK	Slovak Republic
32	SVN	Slovenia
33	SWE	Sweden
34	TUR	Turkey
35	TWN	Taiwan
36	AVE	Sample Average

Table B.2: Sector Names

Sector Index	Sector Code	Sector Name
1	AGR	Agriculture, Hunting, Forestry and Fishing
2	MIN	Mining and Quarrying
3	MAN	Manufacturing
4	UTL	Non-Electrical Utilities and Construction
5	TRD	Trade
6	TRN	Transport
7	HSP	Accommodation and Food Services
8	ICT	Information and Communication Services
9	FIN	Financial Services
10	OTH	Other Private Services
11	SOC	Social and Public Services
	ENG	Energy
	FND	Final Demand

B.2 Derivation of Sectoral Prices

Setting price equal to unit cost, sectoral prices are given by

$$p_{i,t} = A_{i,t}^{-1} w_t^{\omega_{il}} r_{k,t}^{\omega_{ik}} p_{ie,t}^{\omega_{ie}} \left(\prod_{j=1}^{N} p_{j,t}^{\omega_{ij}} \right) \; \forall \; i \in N$$

Let $\underline{p}_t$, $\underline{p}_{e,t}$, $\underline{A}_t$, $\underline{\omega}_l$, $\underline{\omega}_k$ and $\underline{\omega}_e$ denote the $N \times 1$ column vectors with ith elements $p_{i,t}$, $p_{ie,t}$, $A_{i,t}$, ω_{il}, ω_{ik} and ω_{ie} respectively. Taking logs, this implies the log-linear system of sectoral pricing equations given by

$$\begin{aligned}
\ln \underline{p}_t &= -\ln \underline{A}_t + \underline{\omega}_l \ln w_t + \underline{\omega}_k \ln r_{k,t} + \underline{\omega}_e \ln \underline{p}_{e,t} + \Omega \ln \underline{p}_t \\
\Rightarrow \ln \underline{p}_t &= \psi \left[-\ln \underline{A}_t + \underline{\omega}_l \ln w_t + \underline{\omega}_k \ln r_{k,t} + \underline{\omega}_e \ln \underline{p}_{e,t} \right] \\
\Rightarrow p_{i,t} &= \left(\prod_{j=1}^{N} A_{j,t}^{-\psi_{ij}} \right) w_t^{\left(\sum_{j=1}^{N} \psi_{ij} \omega_{jl} \right)} r_{k,t}^{\left(\sum_{j=1}^{N} \psi_{ij} \omega_{jk} \right)} \left(\prod_{j=1}^{N} p_{je,t}^{\psi_{ij} \omega_{je}} \right) \; \forall \; i \in N
\end{aligned}$$

B.3 Derivation of Sales Share Expressions

To derive the sales share expressions, first note that the Cobb-Douglas functional forms and firm FOC imply

$$\begin{aligned} p_{i,t}Y_{i,t} &= \omega_{yi}Y_t \ \forall \ i \in N \\ p_{ye,t}Y_{e,t} &= \omega_{ye}Y_t \end{aligned}$$

$$\begin{aligned} w_t L_{i,t} &= \omega_{il}p_{i,t}Q_{i,t} \ \forall \ i \in N \\ r_{k,t}K_{i,t} &= \omega_{ik}p_{i,t}Q_{i,t} \ \forall \ i \in N \\ p_{ie,t}Z_{ie,t} &= \omega_{ie}p_{i,t}Q_{i,t} \ \forall \ i \in N \\ p_{j,t}Z_{ij,t} &= \omega_{ij}p_{i,t}Q_{i,t} \ \forall \ i,j \in N \end{aligned}$$

Plugging these into the goods market clearing condition for output from sector i

$$\begin{aligned} Q_{i,t} &= Y_{i,t} + \sum_{j=1}^{N} Z_{ji,t} \\ \Rightarrow p_{i,t}Q_{i,t} &= \omega_{yi}Y_t + \sum_{j=1}^{N} \omega_{ji}p_{j,t}Q_{j,t} \end{aligned}$$

Letting $\underline{pQ}_t$ and $\underline{\omega}_y$ denote the $N \times 1$ column vectors with ith elements $p_{i,t}Q_{i,t}$ and ω_{yi}, respectively, this implies

$$\begin{aligned} \underline{pQ}_t &= \psi\underline{\omega}_y Y_t \\ \Rightarrow p_{i,t}Q_{i,t} &= \left(\sum_{j=1}^{N} \omega_{yj}\psi_{ji}\right) Y_t \\ \Rightarrow \lambda_i &= \sum_{j=1}^{N} \omega_{yj}\psi_{ji} \end{aligned}$$

Total expenditure on labour not employed by the energy sector is given by

$$\begin{aligned}
\sum_{i=1}^{N} w_t L_{i,t} =& \sum_{i=1}^{N} \omega_{il} p_{i,t} Q_{i,t} \\
\Rightarrow \sum_{i=1}^{N} w_t L_{i,t} =& \left(\sum_{i=1}^{N} \omega_{il} \lambda_i \right) Y_t \\
\Rightarrow \lambda_l =& \sum_{i=1}^{N} \lambda_i \omega_{il}
\end{aligned}$$

Similarly, aggregate capital expenditure is

$$\begin{aligned}
\sum_{i=1}^{N} r_{k,t} K_{i,t} =& \sum_{i=1}^{N} \omega_{ik} p_{i,t} Q_{i,t} \\
\Rightarrow r_{k,t} K_t =& \left(\sum_{i=1}^{N} \omega_{ik} \lambda_i \right) Y_t \\
\Rightarrow \lambda_k =& \sum_{i=1}^{N} \lambda_i \omega_{ik}
\end{aligned}$$

Aggregate energy expenditures are

$$\begin{aligned}
\sum_{i=1}^{N} p_{ie,t} Z_{ie,t} + p_{ye,t} Y_{e,t} =& \sum_{i=1}^{N} \omega_{ie} p_{i,t} Q_{i,t} + \omega_{ye} Y_t \\
\Rightarrow \sum_{i=1}^{N} p_{ie,t} Z_{ie,t} + p_{ye,t} Y_{e,t} =& \sum_{i \in N_y} \lambda_i \omega_{ie} \\
\Rightarrow \lambda_e =& \sum_{i \in N_y} \lambda_i \omega_{ie}
\end{aligned}$$

Since aggregate GDP is equal to factor income payments and transfers to households

$$\begin{aligned}
Y_t =& w_t L_t + r_{k,t} K_t + \sum_{i \in N_y} \tau_{i,t} Q_{id,t} \\
\Rightarrow Y_t =& \sum_{i=1}^{N} w_t L_{i,t} + r_{k,t} K_t + \sum_{i \in N_y} \left[w_t \left(L_{id,t} + L_{ig,t} \right) + \tau_{i,t} Q_{id,t} \right] \\
\Rightarrow 1 =& \frac{\sum_{i=1}^{N} w_t L_{i,t}}{Y_t} + \frac{r_{k,t} K_t}{Y_t} + \frac{\sum_{i \in N_y} p_{ie,t} Q_{ie,t}}{Y_t} \\
\Rightarrow 1 =& \lambda_l + \lambda_k + \lambda_e
\end{aligned}$$

B.4 Derivation of Quantities

The first order conditions of the composite energy producers imply that

$$p_{id,t}Q_{id,t} = \alpha_{id,t}p_{ie,t}Q_{ie,t} \ \forall \ i \in N_y$$
$$p_{ig,t}Q_{ig,t} = \alpha_{ig,t}p_{ie,t}Q_{ie,t} \ \forall \ i \in N_y$$

where sector i's energy expenditure shares on dirty and green energy are given by

$$\alpha_{id,t} \equiv \frac{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}}}{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}} \ \forall \ i \in N_y$$
$$\alpha_{ig,t} \equiv \frac{\omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}}{\omega_{id}^{\sigma_{ie}} A_{id,t}^{\sigma_{ie}-1} \phi_{i,t}^{1-\sigma_{ie}} + \omega_{ig}^{\sigma_{ie}} A_{ig,t}^{\sigma_{ie}-1}} \ \forall \ i \in N_y$$

Given unit cost pricing by the dirty and green energy producers, it also follows that

$$w_t L_{id,t}\phi_{i,t} = p_{id,t}Q_{id,t} \ \forall \ i \in N_y$$
$$w_t L_{ig,t} = p_{ig,t}Q_{ig,t} \ \forall \ i \in N_y$$

Combining these conditions with the first order conditions of the sectoral producers and the results on equilibrium sales shares implies

$$w_t L_{i,t} = \omega_{il}\lambda_i Y_t \ \forall \ i \in N$$
$$w_t L_{id,t} = \frac{\alpha_{id,t}\lambda_{ie}}{\phi_{i,t}} Y_t \ \forall \ i \in N_y$$
$$w_t L_{ig,t} = \alpha_{ig,t}\lambda_{ie} Y_t \ \forall \ i \in N_y$$

We can thus relate all labour allocations to the allocation of labour to dirty energy production for sector $j \in N_y$

$$L_{i,t} = \frac{\phi_{j,t}\lambda_i\omega_{il}}{\lambda_{je}\alpha_{jd,t}} L_{jd,t} \ \forall \ i \in N$$
$$L_{id,t} = \frac{\phi_{j,t}\frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}}{\lambda_{je}\alpha_{jd,t}} L_{jd,t} \ \forall \ i \in N_y$$
$$L_{ig,t} = \frac{\phi_{j,t}\lambda_{ie}\alpha_{ig,t}}{\lambda_{je}\alpha_{jd,t}} L_{jd,t} \ \forall \ i \in N_y$$

Plugging these expressions into the labour market clearing condition implies

$$L_{id,t} = \left(\frac{\frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}}{\lambda_l + \sum_{j\in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j\in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall\ i \in N_y$$

$$L_{ig,t} = \left(\frac{\lambda_{ie}\alpha_{ig,t}}{\lambda_l + \sum_{j\in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j\in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall\ i \in N_y$$

$$L_{i,t} = \left(\frac{\lambda_i\omega_{il}}{\lambda_l + \sum_{j\in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j\in N_y} \lambda_{je}\alpha_{jg,t}} \right) L_t \ \forall\ i \in N_y$$

The energy allocations are then determined as

$$Q_{id,t} = \left(\frac{\frac{\lambda_{ie}\alpha_{id,t}}{\phi_{i,t}}}{\lambda_l + \sum_{j\in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j\in N_y} \lambda_{je}\alpha_{jg,t}} \right) A_{id,t}L_t \ \forall\ i \in N_y$$

$$Q_{ig,t} = \left(\frac{\lambda_{ie}\alpha_{ig,t}}{\lambda_l + \sum_{j\in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j\in N_y} \lambda_{je}\alpha_{jg,t}} \right) A_{ig,t}L_t \ \forall\ i \in N_y$$

$$Q_{ie,t} = A_{ie,t} \left(\omega_{id} Q_{id,t}^{\frac{\sigma_{ie}-1}{\sigma_{ie}}} + \omega_{ig} Q_{ig,t}^{\frac{\sigma_{ie}-1}{\sigma_{ie}}} \right)^{\frac{\sigma_{ie}}{\sigma_{ie}-1}} \quad \forall\ i \in N_y$$

The size of the aggregate capital stock is then pinned down by the fact that capital earns λ_k fraction of total income and GDP is given by

$$Y_t = w_t L_t + r_{k,t} K_t + \sum_{i\in N_y} \tau_{i,t} Q_{id,t}$$

Combining these two facts along with the expression for $Q_{id,t}$ gives us

$$K_t = \left(\frac{\lambda_k}{1-\lambda_k} \right) \frac{w_t}{r_{k,t}} L_t \left[\frac{\lambda_l + \lambda_e}{\lambda_l + \lambda_e + \sum_{i\in N_y} \left(\frac{1-\phi_{i,t}}{\phi_{i,t}} \right) \lambda_{ie}\alpha_{id,t}} \right]$$

The allocation of the aggregate capital stock across sectors is pinned down by combining each sector's first order condition on capital choice with the capital market clearing condition, which implies

$$K_{i,t} = \frac{\lambda_i\omega_{ik}}{\lambda_k} K_t$$

All other allocations can be solved in a straightforward manner from prices, the value of Y_t and the expenditure share implications of the firm's Cobb-Douglass first order conditions.

B.5 Proof of Proposition 2

Let the set of carbon taxes be denoted by $\boldsymbol{\tau_t} \equiv \{\tau_{i,t}\}_{i \in N_y}$. To sign the derivative $\frac{dQ_{d,t}}{d\tau_{i,t}}$, first note that total emissions can be written as a function of $\boldsymbol{\phi_t} \equiv \{\phi_{i,t}\}_{i \in N_y}$ since

$$Q_{d,t}\left(\boldsymbol{\phi_t}\right) = \sum_{i \in N_y} Q_{id,t} = \left(\frac{\sum_{i \in N_y} A_{id,t} \frac{\lambda_{ie}\alpha_{id,t}(\phi_{i,t})}{\phi_{i,t}}}{\lambda_l + \sum_{i \in N_y} \frac{\lambda_{ie}\alpha_{id,t}(\phi_{i,t})}{\phi_{i,t}} + \sum_{i \in N_y} \lambda_{ie}\alpha_{ig,t}\left(\phi_{i,t}\right)} \right) L_t$$

where I have made explicit the dependence of $\alpha_{id,t}$ and $\alpha_{ig,t}$ on $\phi_{i,t}$. In addition, note that each $\phi_{i,t}$ is a function of $\boldsymbol{\tau_t}$ since $\phi_{i,t} = 1 + \frac{\tau_{i,t}}{w_t(\boldsymbol{\tau_t})}$. From the normalising equation $F_{y,t}\left(w_t, \boldsymbol{\tau_t}\right) = 1$, it can be shown that $\frac{dw_t}{d\tau_{i,t}} < 0$ holding constant all other $\tau_{j,t} \in N_y$ from which it follows that $\frac{d\phi_{j,t}}{d\tau_{i,t}} > 0 \; \forall \; j \in N_y$. Assuming that $\sigma_{ie} > 1$ and $A_{id,t} = A_{d,t} \; \forall \; i \in N_y$, it follows that

$$\begin{aligned}
&\frac{dQ_{d,t}}{d\tau_{i,t}} \propto \frac{d}{d\tau_{i,t}} \left[\frac{\sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}(\phi_{j,t}(\tau_{i,t}))}{\phi_{j,t}(\tau_{i,t})}}{\lambda_l + \sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}(\phi_{j,t}(\tau_{i,t}))}{\phi_{j,t}(\tau_{i,t})} + \sum_{j \in N_y} \lambda_{je}\alpha_{jg,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right)} \right] \\
\Rightarrow &\frac{dQ_{d,t}}{d\tau_{i,t}} \propto \left[\frac{d}{d\tau_{i,t}} \left(\sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right)}{\phi_{j,t}\left(\tau_{i,t}\right)} \right) \right] \left[\lambda_l + \sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} + \sum_{j \in N_y} \lambda_{je}\alpha_{jg,t} \right] \\
&\qquad - \left[\sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} \right] \left[\frac{d}{d\tau_{i,t}} \left(\lambda_l + \sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right)}{\phi_{j,t}\left(\tau_{i,t}\right)} + \sum_{j \in N_y} \lambda_{je}\alpha_{jg,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right) \right) \right] \\
\Rightarrow &\frac{dQ_{d,t}}{d\tau_{i,t}} \propto \left[\frac{d}{d\tau_{i,t}} \left(\sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right)}{\phi_{j,t}\left(\tau_{i,t}\right)} \right) \right] \left[\lambda_l + \sum_{j \in N_y} \lambda_{je}\alpha_{jg,t} \right] \\
&\qquad - \left[\sum_{j \in N_y} \frac{\lambda_{je}\alpha_{jd,t}}{\phi_{j,t}} \right] \left[\frac{d}{d\tau_{i,t}} \left(\sum_{j \in N_y} \lambda_{je}\alpha_{jg,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right) \right) \right] < 0
\end{aligned}$$

since

$$\begin{aligned}
\frac{d}{d\tau_{i,t}} \left(\frac{\alpha_{jd,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right)}{\phi_{j,t}\left(\tau_{i,t}\right)} \right) &< 0 \\
\frac{d}{d\tau_{i,t}} \alpha_{jg,t}\left(\phi_{j,t}\left(\tau_{i,t}\right)\right) &> 0
\end{aligned}$$

which follow from

$$\frac{d\phi_{j,t}}{d\tau_{i,t}} \geq 0$$
$$\frac{d\alpha_{jd,t}}{d\phi_{j,t}} < 0$$
$$\frac{d\alpha_{jg,t}}{d\phi_{j,t}} > 0$$

Note that one implication of these results is that raising $\tau_{i,t}$ to cause a shift in $\alpha_{ig,t}$ also raises $\alpha_{jg,t}$ in any sector with $\tau_{j,t} > 0$ since the tax-to-wage ratio increases in sector j (note that this effect is absent when $\tau_{j,t} = 0$ and also note that the taxes are all specified in real terms here such that raising $\tau_{i,t}$ in real terms while holding $\tau_{j,t}$ fixed in real terms would mean nominal tax changes when the price level is changing). Letting N_0 denote the subset of N_y with zero carbon taxes, it follows that

$$\lim_{\tau_{i,t}\to\infty} \phi_{j,t} = \infty \text{ if } j \notin N_0$$
$$\lim_{\tau_{i,t}\to\infty} \phi_{j,t} = 1 \text{ if } j \in N_0$$

from which it follows that

$$\lim_{\tau_{i,t}\to\infty} Q_{d,t} = \left(\frac{\sum_{j\in N_0} A_{jd,t}\lambda_{je}\alpha_{jd,t}}{\lambda_l + \lambda_e} \right) L_t$$

B.6 Proof of Proposition 3

Given the small open economy assumption, maximising individual welfare over the transition period is equivalent to maximising the present discounted value of GDP over this period. The policymaker's problem can therefore be written as

$$\max_{\left\{\{\tau_{i,t}\}_{i\in N_y}\right\}_{t=1}^{T}} \quad \sum_{t=1}^{\infty} \frac{Y_t\left(\{\tau_{i,t}\}_{i\in N_y}\right)}{R_t}$$

$$\text{s.t. } \sum_{t=1}^{T}\sum_{i\in N_y} \bar{Q}_{id,t} \leq \bar{Q}_d$$

where I define

$$R_t \equiv \prod_{s=2}^{t}\left(1+r_{b,s}\right)$$

To build intuition, it helps to start tackling this problem by deriving the first order conditions for the optimal path of the quantity of emissions before translating these into the first order conditions for the optimal carbon tax.

To solve this problem in a compact way, note that we can express real GDP (Y_t) as a reduced form function of the sectoral emissions constraints (where I let $\bar{Q}_{id,t}$ denote the quantity of emissions permits allocated to sector i in period t)[1]:

$$Y_t = G_t\left(\left\{\bar{Q}_{id,t}\right\}_{i\in N_y}\right)$$

It follows that the planner's optimal policy problem can then be written as

$$\max_{\left\{\left\{\bar{Q}_{id,t}\right\}_{i\in N_y}\right\}_{t=1}^{T}} \quad \sum_{t=1}^{\infty} \frac{G_t\left(\left\{\bar{Q}_{id,t}\right\}_{i\in N_y}\right)}{R_t}$$

$$\text{s.t. } \sum_{t=1}^{T}\sum_{i\in N_y} \bar{Q}_{id,t} \leq \bar{Q}_d$$

Assuming a binding constraint on emissions, it follows that the optimal solution satisfies the

[1] As shown in Subsections 2.3.2 and 2.3.3, the reduced form production function always exists and is unique.

following conditions

$$\frac{dG_t}{d\bar{Q}_{id,t}} = \frac{dG_t}{d\bar{Q}_{jd,t}} \ \forall\ t \text{ and } i,j \in N_y$$

$$\frac{dG_t\left(\bar{Q}_{id,t}\right)}{d\bar{Q}_{id,t}} = \frac{1}{1+r_{t+1}}\frac{dG_{t+1}\left(\bar{Q}_{id,t+1}\right)}{d\bar{Q}_{id,t+1}} \ \forall\ t \text{ and } i$$

$$\sum_{t=1}^{T}\sum_{i\in N_y}\bar{Q}_{id,t} = \bar{Q}_d$$

The first condition is a static efficiency condition: within a given period, emissions should be allocated across sectors such that their marginal impact on GDP is equalised. The second condition is a dynamic efficiency condition: across periods, the marginal impact of emissions on GDP should grow at a rate equal to the real interest rate. We can then use the envelope theorem to relate these conditions on quantities to an equivalent set of conditions on prices. Letting variables with stars denote the static competitive equilibrium outcomes conditional on policy choice $\left\{\bar{Q}_{id,t}\right\}_{i\in N_y}$, factor endowment L_t, and capital rental rate $r_{k,t}$, we can write the aggregate production function as

$$\begin{aligned}
&G_t\left(L_t, r_{k,t}, \left\{\bar{Q}_{id,t}\right\}_{i\in N_y}\right)\\
&= \left(\frac{Y^*_{e,t}}{\omega_{ye}}\right)^{\omega_{ye}}\left[\prod_{i=1}^{N}\left(\frac{Y^*_{i,t}}{\omega_{yi}}\right)^{\omega_{yi}}\right]\\
&-\sum_{i=1}^{N}p^*_{i,t}\left\{Y^*_{i,t}+\sum_{j=1}^{N}Z^*_{ji,t}-A_{i,t}\left(\frac{L^*_{i,t}}{\omega_{il}}\right)^{\omega_{il}}\left(\frac{K^*_{i,t}}{\omega_{ik}}\right)^{\omega_{ik}}\left(\frac{Z^*_{ie,t}}{\omega_{ie}}\right)^{\omega_{ie}}\left[\prod_{j=1}^{N}\left(\frac{Z^*_{ij,t}}{\omega_{ij}}\right)^{\omega_{ij}}\right]\right\}\\
&-\sum_{i=1}^{N}p^*_{ie,t}\left\{Z^*_{ie,t}-A_{ie,t}\left(\omega_{id}Q^{*\frac{\sigma_{ie}-1}{\sigma_{ie}}}_{id,t}+\omega_{ig}Q^{*\frac{\sigma_{ie}-1}{\sigma_{ie}}}_{ig,t}\right)^{\frac{\sigma_{ie}}{\sigma_{ie}-1}}\right\}\\
&-p^*_{ye,t}\left\{Y^*_{e,t}-A_{ye,t}\left(\omega_{yd}Q^{*\frac{\sigma_{ye}-1}{\sigma_{ye}}}_{yd,t}+\omega_{yg}Q^{*\frac{\sigma_{ye}-1}{\sigma_{ye}}}_{yg,t}\right)^{\frac{\sigma_{ye}}{\sigma_{ye}-1}}\right\}\\
&-\sum_{k\in\{d,g\}}\sum_{i\in N_y}p^*_{ik,t}\left\{Q^*_{ik,t}-A_{ik,t}L^*_{ik,t}\right\}-w^*_t\left\{\sum_{i=1}^{N}L^*_{i,t}+\sum_{k\in\{d,g\}}\sum_{i\in N_y}L^*_{ik,t}-L_t\right\}\\
&-r_{k,t}\left\{\sum_{i=1}^{N}K^*_{i,t}-K^*_t\right\}-\sum_{i\in N_y}\tau^*_{it}\left\{Q^*_{id,t}-\bar{Q}_{id,t}\right\}
\end{aligned}$$

where the fact that all competitive equilibrium variables are functions of $\left\{\bar{Q}_{id,t}\right\}_{i\in N_y}$, L_t, and $r_{k,t}$ has not been made explicit for notational convenience. Differentiating both sides of this expression by $\bar{Q}_{id,t}$ and plugging in the competitive equilibrium conditions from Subsections

2.3.2 and 2.3.3 leaves us with

$$\frac{dG_t\left(L_t, r_{k,t}, \left\{\bar{Q}_{id,t}\right\}_{i\in N_y}\right)}{d\bar{Q}_{id,t}} = \tau^*_{it}$$

This result follows from the standard envelope theorem intuition: the increase in GDP associated with a marginal increase in the quantity of emissions allocated to sector i is equal to the price put on emissions in that sector. Combining this result with the characterisation of the optimal quantities implies that the optimal carbon tax policy must satisfy the following conditions

$$\tau_{i,t} = \tau_{j,t} \; \forall \; t \text{ and } i,j \in N_y$$
$$\tau_{i,t} = \frac{1}{1+r_{t+1}}\tau_{i,t+1} \; \forall \; t \text{ and } i$$
$$\sum_{t=1}^{T}\sum_{i\in N_y} Q_{id,t}\left(\tau_{i,t}\right) = \bar{Q}_d$$

The first condition states that within a period, carbon prices should be equalised across sectors. The second condition gives us the Hotelling rule: the price of emissions (a scarce and exhaustible resource) should optimally grow at a rate equal to the real interest rate. To show that the solution to these conditions exists and is unique, define the excess emissions function as

$$H\left(\left\{\left\{\tau_{i,t}\right\}_{i\in N_y}\right\}_{t=1}^{T}\right) \equiv \sum_{t=1}^{T}\sum_{i\in N_y} Q^*_{id,t}\left(\tau_{i,t}\right) - \bar{Q}_d$$

It follows that under the optimal carbon tax policy, this excess emissions can be expressed solely as a function of the period 1 carbon tax common across all sectors:

$$H\left(\tau_1\right) = \sum_{t=1}^{T} Q_{d,t}\left(R_t\tau_1\right) - \bar{Q}_d$$

A direct corollary of Proposition 2 is that

$$\frac{dH\left(\tau_1\right)}{d\tau_1} < 0$$

and

$$\lim_{\tau_1\to\infty} H\left(\tau_1\right) < 0$$

Furthermore, given the assumption that the carbon budget is a binding constraint, $H\left(0\right) > 0$.

It thus follows that a unique τ_1^* exists such that $H(\tau_1^*) = 0$. Given the uniqueness and existence of τ_1^*, it follows from the condition that $\tau_t^* = \frac{\tau_1^*}{R_t}$ that the sequence $\{\tau_t^*\}_{t=1}^{T}$ also exists and is unique.

B.7 Suboptimal Policy Problem

Let N_0 denote the subset of all sectors that face no carbon tax and N_τ denote the subset of all sectors that face a common carbon tax $\tau_t > 0$. Furthermore, define the sales shares of exempt and non-exempt energy as

$$\lambda_{0e} = \sum_{i \in N_0} \lambda_{ie}$$
$$\lambda_{\tau e} = \sum_{i \in N_\tau} \lambda_{ie}$$

Equation (2.36) then becomes

$$F_{p,t}(w_t) \equiv A_t^{-1} w_t^{1-\lambda_k} r_{k,t}^{\lambda_k} \left[\left(\omega_d^{\sigma_e} A_{d,t}^{\sigma_e - 1} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e - 1} \right)^{\frac{\lambda_{0e}}{\lambda_e}} \left(\omega_d^{\sigma_e} A_{d,t}^{\sigma_e - 1} \phi_t^{1-\sigma_e} + \omega_g^{\sigma_e} A_{g,t}^{\sigma_e - 1} \right)^{\frac{1-\lambda_{0e}}{\lambda_e}} \right]^{\frac{\lambda_e}{1-\sigma_e}}$$

As a special case of Proposition 1, this equation implies that w_t is a function of τ_t and $\hat{\lambda}_{0e} \equiv \frac{\lambda_{0e}}{\lambda_e}$. Equation (2.44) then implies that $Q_{d,t}$ is also a function of τ_t and $\hat{\lambda}_{0e}$. The sub-optimal policy problem can thus be written as

$$\max_{\{\tau_t\}_{t=1}^{T}} \quad \sum_{t=1}^{\infty} \frac{Y_t\left(\tau_t, \hat{\lambda}_{0e}\right)}{R_t}$$

$$\text{s.t.} \sum_{t=1}^{T} Q_{d,t}\left(\tau_t, \hat{\lambda}_{0e}\right) \leq \bar{Q}_d$$

Using the results established in the proof of Proposition 3 (see Appendix B.6), it follows that intertemporal optimisation still requires that the sectorally suboptimal carbon tax grows over time at a rate equal to the real interest rate. Because $\frac{dQ_{d,t}}{d\hat{\lambda}_{0e}} > 0$, it follows that $\frac{d\tau_t^*}{d\hat{\lambda}_{0e}} > 0$.

B.8 Additional Tables

Table B.3: Optimal Policy by Country: Carbon Tax and Discounted Output Losses

Country	τ^*_{2020}	$\%\Delta Y$	$\frac{Q_{0d}}{Q_d}$	p_{0d}
TWN	749	-5.17	3.7	47
KOR	713	-6.23	4.1	38
BEL	674	-4.37	3.3	54
CZE	585	-4.92	3.4	40
AUS	545	-5.47	7.2	18
LUX	533	-3.71	6.9	18
NLD	496	-2.90	3.3	39
POL	490	-4.06	2.8	41
EST	473	-4.20	3.3	33
CAN	436	-4.48	6.5	17
FIN	435	-2.89	3.4	34
SVK	410	-2.53	2.0	56
BGR	399	-2.28	1.9	59
NOR	340	-1.86	3.9	27
AUT	328	-1.75	2.9	32
DNK	320	-1.06	2.2	48
LTU	304	-0.69	1.4	100
IRL	295	-1.27	3.5	25
GBR	288	-1.06	2.3	36
HUN	286	-0.96	1.6	62
SWE	270	-0.77	1.9	53
SVN	263	-1.49	2.5	28
ESP	260	-0.96	1.9	48
ITA	256	-1.11	2.0	41
FRA	249	-0.83	1.9	45
CHE	243	-0.42	1.8	55
GRC	240	-1.57	2.1	31
CYP	216	-1.86	2.5	21
TUR	201	-1.31	1.9	33
PRT	190	-0.57	1.6	43
HRV	141	-0.37	1.5	36
ROU	140	-0.21	1.2	84
LVA	119	-0.19	1.2	64
MEX	89	-0.23	1.3	33
MLT	4	0.00	1.0	100
AVE	325	-2.1	2.6	42

Notes: Column two of the table provides the value of optimal carbon tax in 2020 measured in 2017 PPP \$ (τ^*_{2020}). Column three is the reduction in the present value of real GDP between 2020 and 2049 from implementing the optimal policy ($\%\Delta Y$). Column four is the ratio of emissions with a zero carbon tax to each country's emissions budget from 2020-2049 ($\frac{Q_{0d}}{Q_d}$). Column five is the 2020 value for the unregulated dirty energy price indexed to the maximum price in the sample (p_{0d}). All calculations are done under the baseline calibration values.

Table B.4: Contribution to λ_e by Sector (Percentage)

	Sector Code											
Country Code	AGR	MIN	MAN	UTL	TRD	TRN	HSP	ICT	FIN	OTH	SOC	FND
AUS	3.5	10.6	15.6	7.3	3.2	16.7	1.6	1.8	0.3	3.1	4.8	31.4
AUT	2.5	1.1	25.8	3.8	4.0	10.7	2.0	0.8	0.7	3.6	8.0	36.9
BEL	2.1	0.7	21.2	2.0	3.6	7.6	0.6	1.0	0.6	3.5	3.3	53.8
BGR	2.9	5.6	19.5	4.2	9.3	14.0	0.7	0.6	0.0	2.8	3.6	36.9
CAN	3.6	3.7	19.8	4.2	4.7	9.7	0.9	0.6	1.4	4.9	7.1	39.4
CHE	1.8	1.1	24.4	2.1	4.3	7.4	1.5	0.9	1.0	5.8	13.9	36.0
CYP	7.3	1.5	20.3	4.6	5.9	13.6	5.0	3.4	4.0	3.0	6.1	25.3
CZE	2.3	1.4	29.7	3.4	6.9	8.7	0.9	0.8	1.0	5.0	5.1	34.9
DNK	3.7	0.2	8.9	2.6	3.4	37.2	0.7	0.8	0.3	1.7	4.2	36.2
ESP	1.5	1.3	27.2	3.5	8.8	10.3	1.8	1.5	0.7	2.5	4.8	36.2
EST	6.2	1.2	18.4	5.5	5.0	21.9	1.1	0.8	0.3	4.6	7.6	27.7
FIN	2.6	1.3	23.8	3.3	3.0	13.3	0.6	0.8	0.8	12.6	6.0	31.9
FRA	2.8	0.4	21.2	3.1	6.3	9.4	0.9	2.8	0.6	3.8	5.4	43.4
GBR	1.9	0.8	22.0	2.8	5.5	6.0	1.6	1.7	1.8	3.5	5.7	46.7
GRC	4.8	0.1	9.2	2.1	3.5	7.7	1.4	0.9	0.3	1.6	3.2	65.2
HRV	2.8	2.2	15.6	5.3	9.4	8.0	3.9	2.8	1.4	6.5	9.3	32.7
HUN	4.1	1.0	32.4	2.8	5.6	10.9	0.8	1.0	0.5	3.6	4.6	32.6
IRL	1.8	3.8	20.6	4.6	8.1	4.1	1.0	6.1	3.3	11.1	9.3	26.1
ITA	2.4	0.5	24.5	4.5	8.4	9.1	1.9	0.8	0.6	3.3	6.6	37.4
KOR	1.4	0.8	37.0	3.3	3.6	12.1	1.2	0.9	0.7	3.7	6.5	28.8
LTU	4.5	0.4	8.2	1.9	1.6	7.2	0.4	0.3	0.3	1.9	3.9	69.5
LUX	1.1	0.1	31.4	2.4	4.0	16.2	1.6	1.9	11.6	9.3	5.7	14.7
LVA	9.2	0.5	18.2	4.8	6.8	20.7	1.0	1.2	0.5	7.2	7.3	22.5
MEX	2.2	1.7	21.5	3.3	5.6	19.9	2.4	0.7	0.4	4.3	5.7	32.4
MLT	2.5	0.3	19.1	3.2	7.6	21.8	4.6	0.8	1.5	5.9	2.5	30.2
NLD	3.4	1.8	18.1	2.3	4.1	10.3	1.1	0.5	0.6	2.8	4.0	51.1
NOR	3.1	2.7	14.1	2.7	3.5	11.0	0.7	0.7	0.3	2.9	7.7	50.6
POL	4.0	1.6	21.2	5.1	6.5	12.4	0.7	0.7	1.4	11.3	4.9	30.2
PRT	3.5	0.9	20.5	5.4	5.5	12.2	2.5	0.8	0.6	2.2	6.4	39.4
ROU	5.6	1.6	29.2	8.0	7.0	6.4	1.4	1.4	1.0	5.6	2.3	30.5
SVK	3.3	0.8	27.0	2.1	4.1	10.0	0.8	0.9	0.2	2.9	5.0	42.9
SVN	2.9	1.0	33.0	4.9	6.6	10.5	2.1	1.4	0.9	3.9	5.4	27.5
SWE	2.4	1.1	21.6	3.3	2.6	11.9	0.6	0.8	0.3	6.1	5.0	44.3
TUR	5.5	2.4	28.1	4.3	6.5	17.2	2.3	0.9	1.7	2.8	6.2	22.0
TWN	2.4	1.4	48.5	1.5	5.8	12.3	1.4	0.4	0.3	1.1	3.1	21.7
AVE	3.3	1.6	22.8	3.7	5.4	12.5	1.5	1.3	1.2	4.6	5.7	36.3

Notes: This table shows each sector's contribution to overall energy expenditures in the economy as a percentage of total energy spending.

Table B.5: Standard Deviation of Sales/Factor Shares (Percentage Points, 2000-2014)

	Sectors												Factors	
	AGR	MIN	MAN	UTL	TRD	TRN	HSP	ICT	FIN	OTH	SOC	ENG	LAB	CAP
AUS	1.3	3.4	5.8	2.5	1.4	0.5	0.4	0.7	0.9	1.3	0.7	0.4	1.1	1.3
AUT	0.2	0.3	2.6	1.0	0.4	0.7	0.2	0.3	0.2	1.8	0.7	1.0	1.4	1.0
BEL	0.4	0.4	5.5	2.3	2.2	0.7	0.2	0.7	0.8	1.7	1.1	1.6	1.3	1.1
BGR	4.2	1.3	5.2	5.1	1.4	2.5	0.1	0.8	1.9	2.0	2.1	1.3	3.1	3.1
CAN	0.5	1.5	5.0	1.4	0.5	0.2	0.2	0.2	0.2	0.6	0.7	0.4	1.0	0.7
CHE	0.3	0.1	2.3	0.4	0.6	0.7	0.4	0.4	1.2	1.4	1.0	0.2	1.9	1.9
CYP	0.8	0.1	2.9	3.2	0.8	4.1	1.0	1.1	2.3	1.7	1.4	1.1	1.6	1.6
CZE	0.9	0.3	7.5	1.5	1.3	0.4	0.4	0.3	0.3	1.7	1.2	0.9	1.0	1.6
DNK	0.6	0.3	2.8	0.7	0.8	2.2	0.1	0.4	0.9	2.0	1.1	1.4	1.4	1.7
ESP	0.8	0.2	4.5	7.5	0.5	0.3	0.3	0.3	0.5	1.4	2.1	1.4	1.6	0.8
EST	1.3	0.2	6.9	2.2	1.2	1.9	0.1	0.5	0.5	1.2	1.4	1.2	2.4	2.8
FIN	0.3	0.4	6.3	0.8	0.5	0.5	0.1	0.7	0.3	2.8	2.3	1.6	1.6	2.9
FRA	0.4	0.1	4.4	0.7	0.4	0.2	0.2	0.2	0.7	1.3	0.8	0.8	1.0	1.4
GBR	0.1	0.2	3.2	0.8	1.5	0.5	0.3	0.4	1.4	1.7	1.6	1.0	1.1	0.8
GRC	1.3	0.2	2.5	4.4	2.2	1.2	0.7	0.5	0.4	1.4	1.2	2.5	2.5	2.7
HRV	1.5	0.8	3.2	2.8	1.9	0.8	0.5	0.2	0.6	2.1	1.3	1.0	3.4	2.6
HUN	2.0	0.3	5.8	1.3	0.4	1.0	0.2	0.5	0.5	1.1	1.2	1.5	2.0	1.2
IRL	0.9	0.4	10.3	6.0	1.4	1.0	1.1	3.2	6.0	4.5	1.5	0.4	2.5	2.7
ITA	0.3	0.3	3.0	1.0	1.0	0.2	0.1	0.6	0.5	0.7	1.1	0.8	1.1	1.6
KOR	0.8	0.7	9.1	0.9	0.6	0.5	0.2	0.7	0.6	0.7	1.0	2.6	2.6	0.8
LTU	1.0	0.2	1.8	2.6	0.9	1.9	0.2	1.1	0.5	1.2	2.6	2.2	2.9	2.0
LUX	0.2	0.1	4.8	0.5	4.1	0.9	0.4	2.9	18.9	8.8	2.3	0.4	2.1	2.1
LVA	0.9	0.2	3.8	4.7	1.0	3.0	0.3	0.7	0.5	3.8	1.3	1.4	2.9	3.2
MEX	0.3	0.7	3.1	0.5	0.6	0.4	0.5	0.2	0.6	0.9	0.6	0.9	2.0	1.3
MLT	0.3	0.2	8.1	1.1	0.8	5.3	0.8	2.1	12.0	5.8	4.9	1.1	1.4	1.7
NLD	0.9	0.8	4.5	1.4	0.9	0.3	0.2	0.5	0.5	0.5	1.1	1.4	1.8	1.1
NOR	0.5	2.1	2.3	1.7	0.5	0.9	0.2	0.5	0.7	1.4	1.3	0.4	1.9	1.8
POL	0.6	0.3	5.4	1.8	2.5	0.9	0.1	0.5	0.3	1.0	0.9	1.1	2.8	2.0
PRT	0.5	0.2	4.8	3.8	1.0	0.7	0.1	0.5	0.7	0.8	1.0	1.5	2.8	1.5
ROU	5.0	0.6	2.4	4.9	3.2	1.7	0.5	0.8	0.8	2.6	1.2	1.0	2.1	2.0
SVK	1.5	0.2	11.1	1.6	1.2	2.5	0.2	0.3	0.4	1.1	1.3	0.9	1.2	1.0
SVN	0.7	0.1	3.7	2.5	0.9	1.1	0.2	0.5	0.4	1.0	1.1	0.7	1.4	1.4
SWE	0.2	0.3	6.0	0.8	0.9	0.5	0.2	0.2	0.3	1.0	0.9	1.1	1.1	1.5
TUR	1.5	0.6	3.4	0.6	0.9	0.7	0.2	0.3	2.7	2.2	0.5	0.6	1.1	0.8
TWN	0.3	0.4	11.0	0.2	0.5	0.4	0.3	0.4	0.9	0.4	1.6	2.5	2.7	0.6
AVE	1.0	0.5	5.0	2.1	1.2	1.2	0.3	0.7	1.7	1.9	1.4	1.1	1.9	1.7
SD	1.0	0.6	2.5	1.8	0.8	1.1	0.2	0.7	3.6	1.6	0.8	0.6	0.7	0.7

Notes: This table shows the standard deviation in sales and factor shares for each country between 2000 and 2014. The LAB and CAP columns refer to the labour and capital shares, respectively (excluding value added in the energy sector). The AVE and SD rows are the average and standard deviations across all countries, respectively.

Appendix C

Appendices to Chapter 3

C.1 Country and Sector Aggregation Scheme

Table C.1: Country/Region Aggregation Scheme

Index	Code	Name	Constituent Country Codes (for Aggregate Regions)
1	BGR	Bulgaria	N/A
2	BRA	Brazil	N/A
3	CHN	China	N/A
4	IDN	Indonesia	N/A
5	IND	India	N/A
6	MEX	Mexico	N/A
7	ROU	Romania	N/A
8	RUS	Russia	N/A
9	TUR	Turkey	N/A
10	EAP	East Asia and Pacific	AUS, JPN, KOR, TWN
11	EUR	Europe	AUT, BEL, CHE, CYP, CZE, DEU, DNK, ESP, EST, FIN, FRA, GBR, GRC, HRV, HUN, IRL, ITA, LTU, LUX, LVA, MLT, NLD, NOR, POL, PRT, SVK, SVN, SWE
12	NAC	North America	CAN, USA
13	ROW	Rest of the World	(Countries not mentioned above)
14	WLD	World	(All countries)

Notes: Country codes are consistent with those used in the 2016 version of the WIOD.

Table C.2: Sector Aggregation Scheme

Sector Index	Sector Name	Corresponding WIOD 2016 Sector
1	Agriculture and Mining	1-4
2	Manufacturing	5-22
3	Services	23-56

Notes: Sector numbers are consistent with those used in the 2016 version of the WIOD.

C.2 Figure 3.9 and Table 3.3: Missing Data

Figure C.1: Cross-Country Distribution of Missing Trade Balance Data (1995-2019)

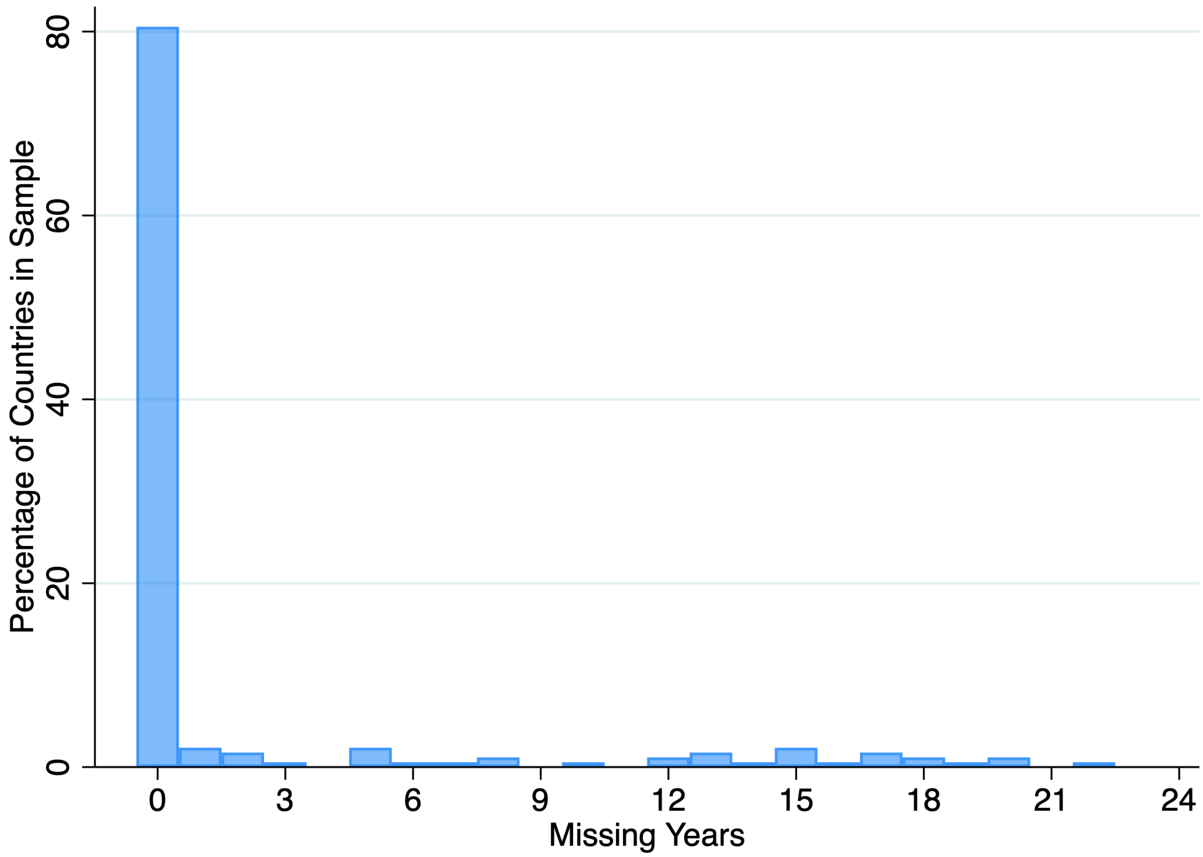

Notes: This figure illustrates the cross-country distribution of missing trade balance data between 1995 and 2019.
Source: World Bank (2022); Author's calculations.

Figure C.2: Cross-Country Distribution of Missing Regression Data (1995-2019)

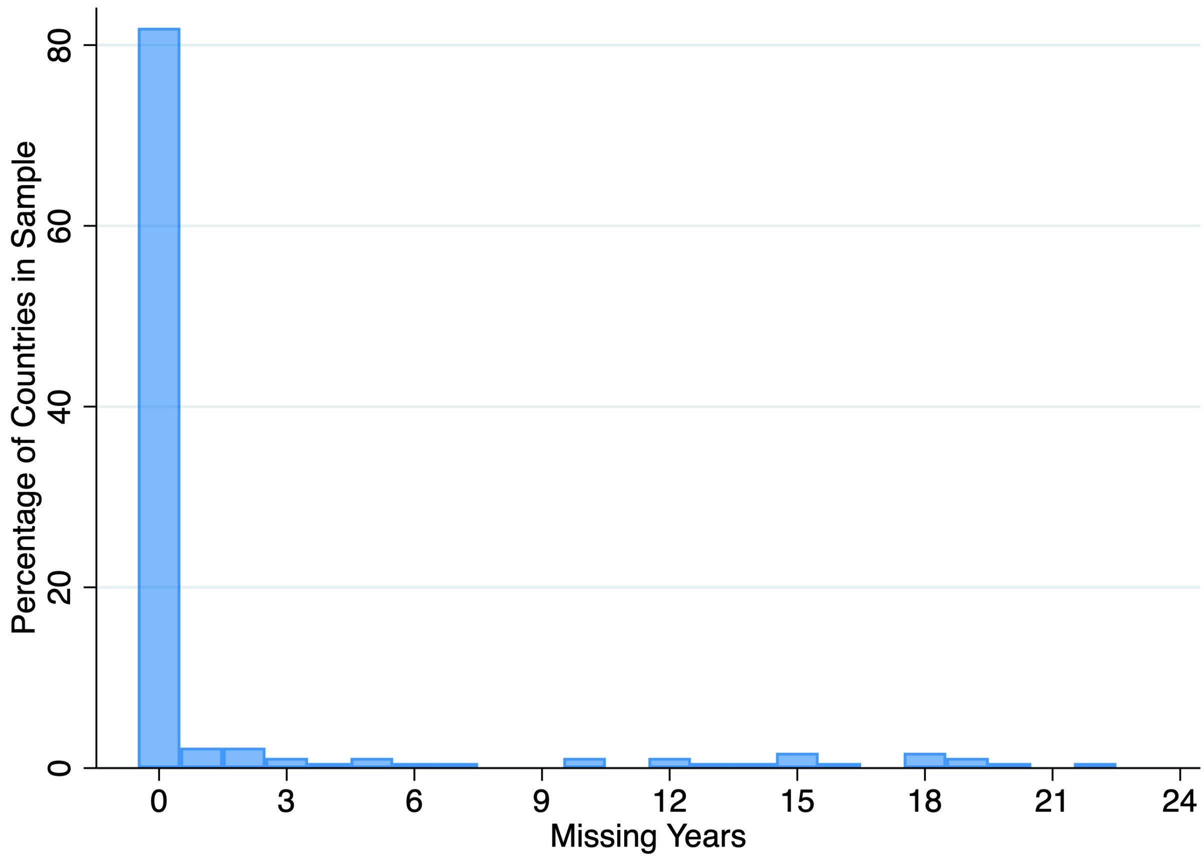

Notes: This figure illustrates the cross-country distribution of missing observations between 1995 and 2019 from regression (3.2).
Source: World Bank (2022); Author's calculations.

C.3 Figure 3.8: Missing Data

Figure C.3: Cross-Country Distribution of Missing Tourism Receipts Data (1995-2019)

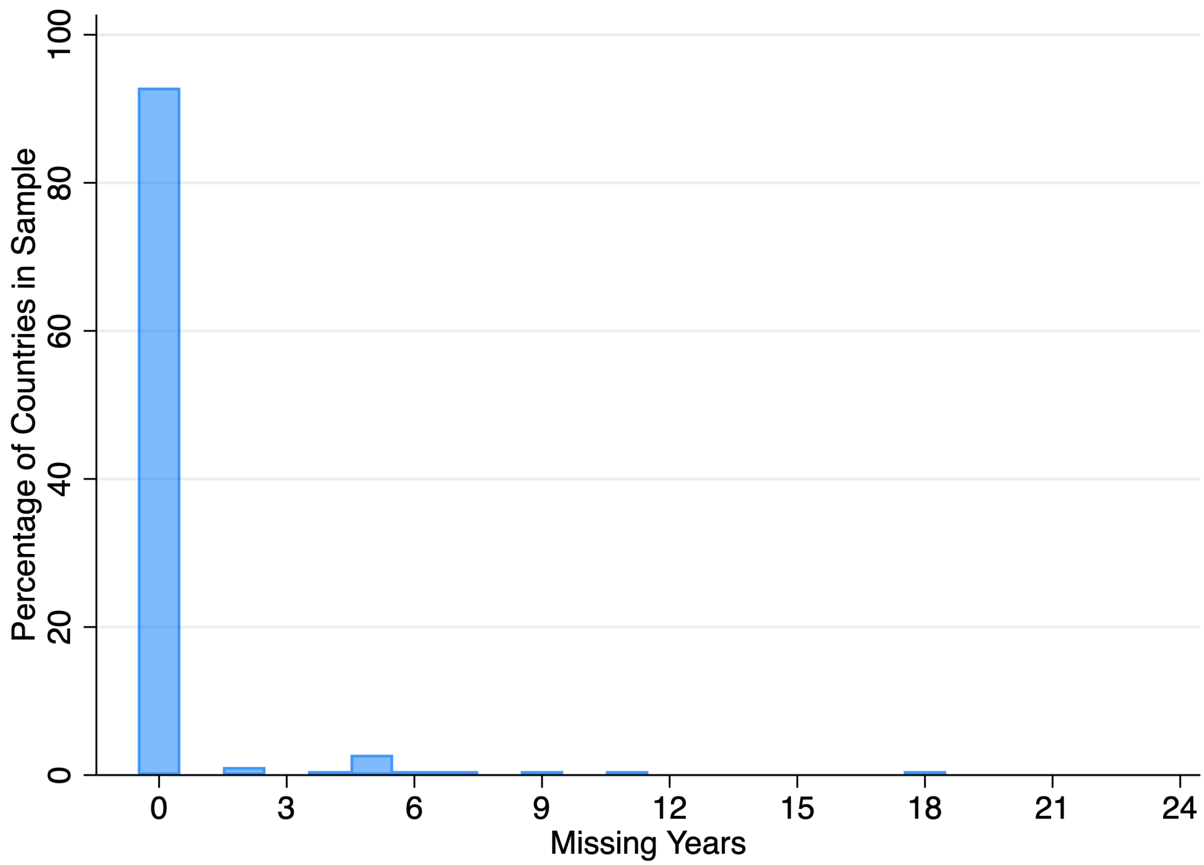

Notes: This figure illustrates the cross-country distribution of missing observations between 1995 and 2019 from Figure 3.8.
Source: World Bank (2022); Author's calculations.

C.4 Baseline Data (2000-2014)

Figure C.4: Nominal Value Added by Sector (% of GDP)

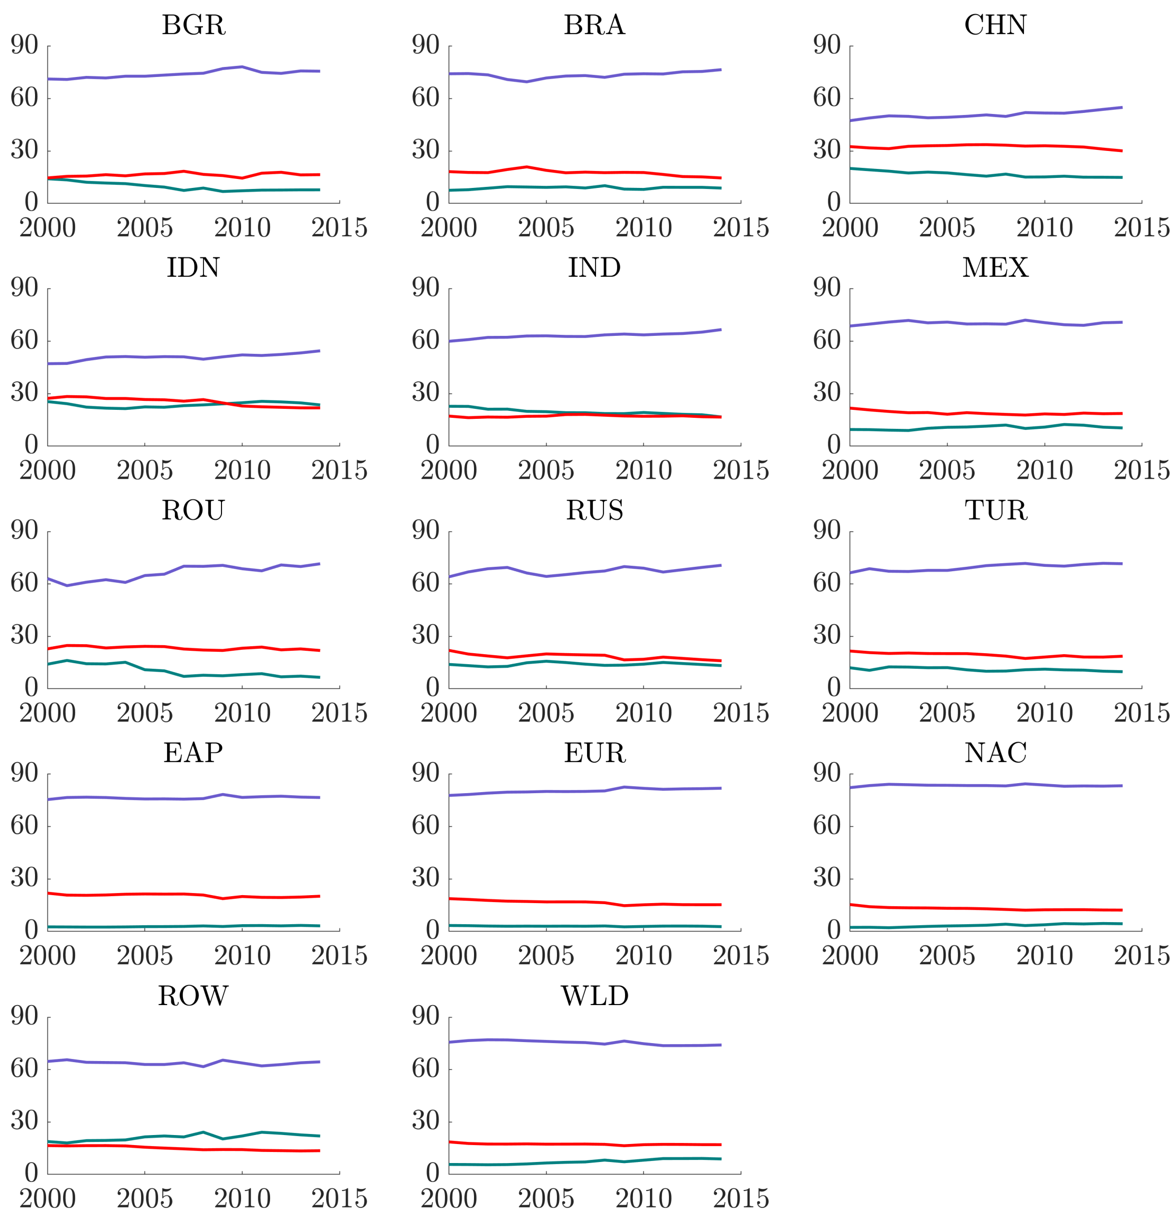

Notes: This figure illustrates nominal value added by sector as a percentage of GDP from 2000 to 2014 for each country/region in the model. The green, red and purple lines represent the share of the Agriculture, Manufacturing and Service sectors, respectively.
Source: Timmer et al. (2015); Author's calculations.

Figure C.5: Real Value Added by Sector (% of GDP)

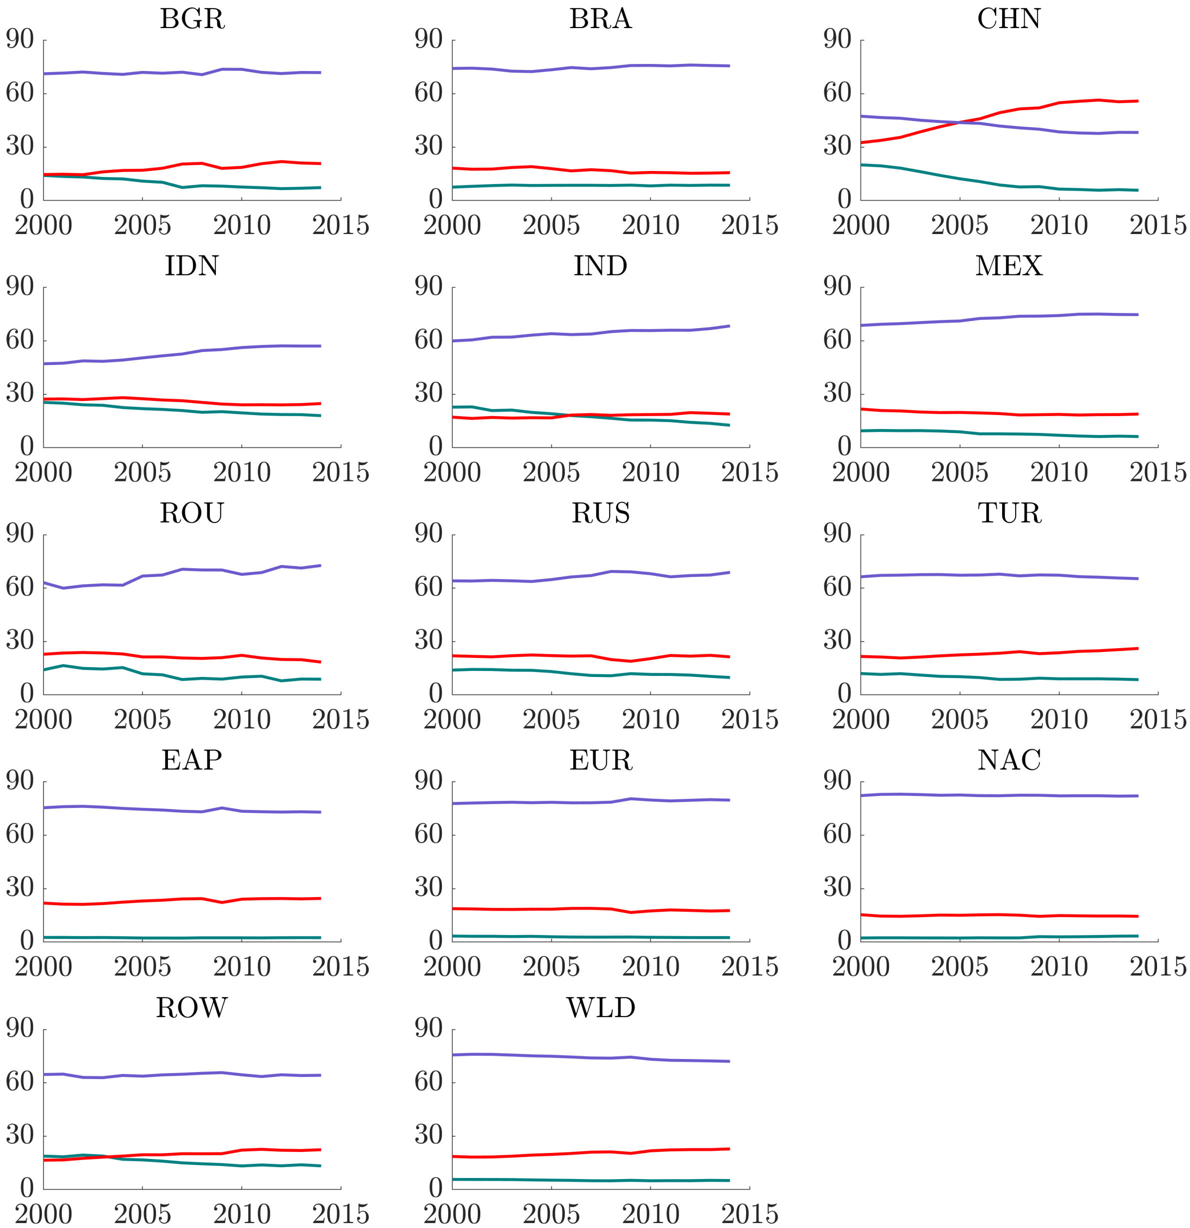

Notes: This figure illustrates real value added by sector as a percentage of GDP from 2000 to 2014 for each country/region in the model. The green, red and purple lines represent the share of the Agriculture, Manufacturing and Service sectors, respectively. See Tables C.1 and C.2 in Appendix C.1 for details on country/region codes and sectoral aggregation.
Source: Timmer et al. (2015); Author's calculations.

Figure C.6: Sectoral Employment Shares

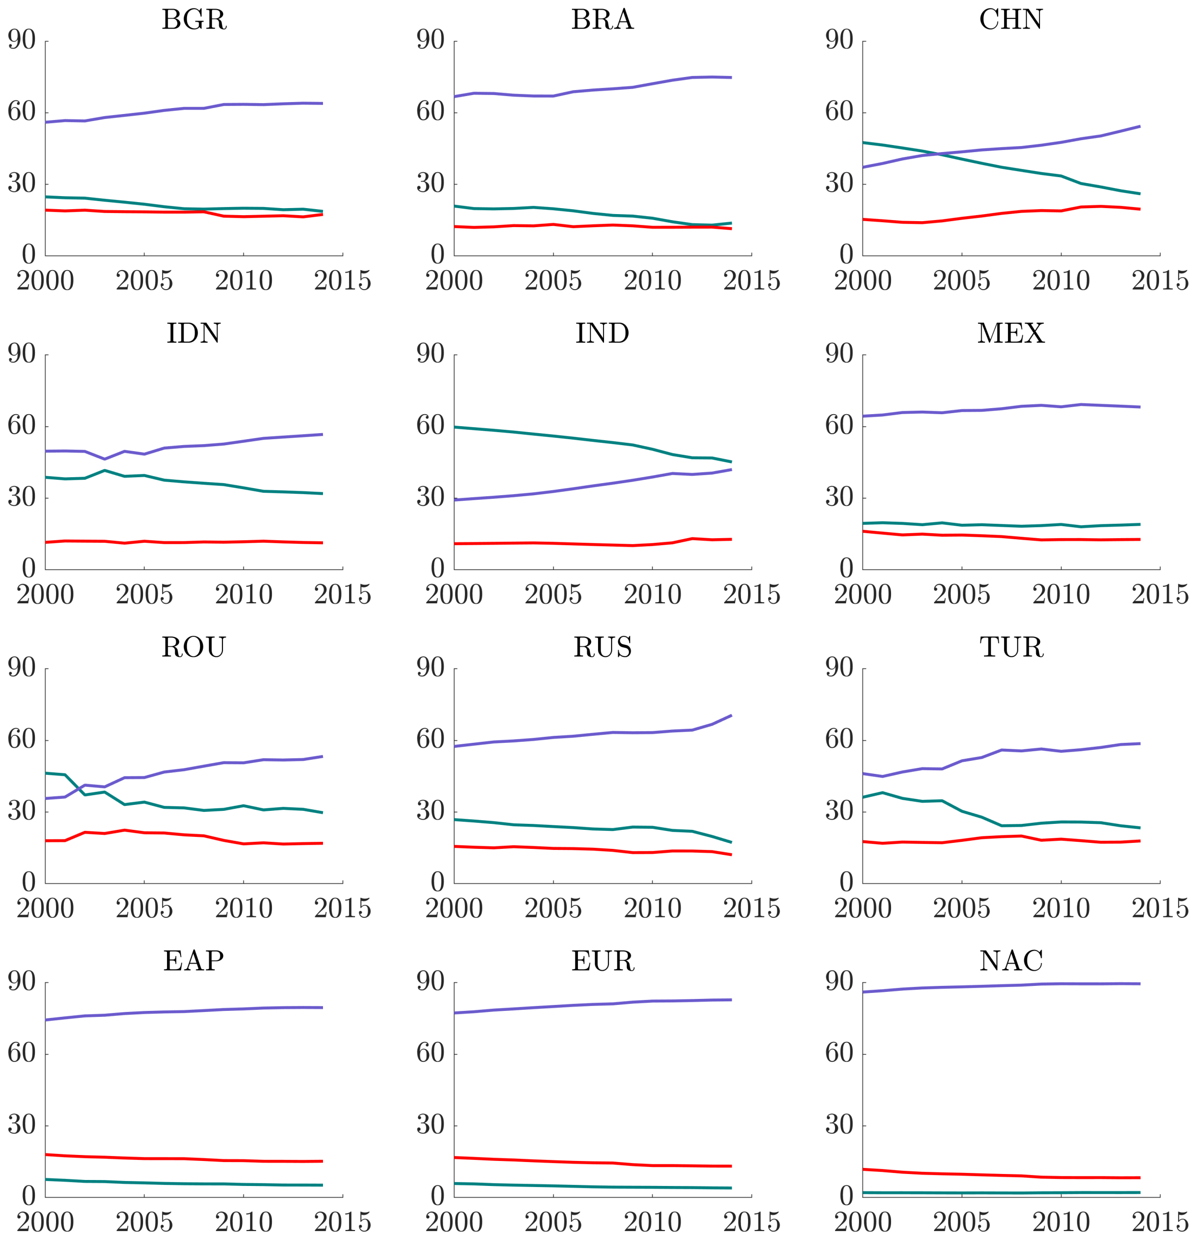

Notes: This figure illustrates the sectoral employment shares from 2000 to 2014 for each country/region in the model. The green, red and purple lines represent the share of the Agriculture, Manufacturing and Service sectors, respectively. See Tables C.1 and C.2 in Appendix C.1 for details on country/region codes and sectoral aggregation.
Source: Timmer et al. (2015); Author's calculations.

Figure C.7: Manufacturing Share

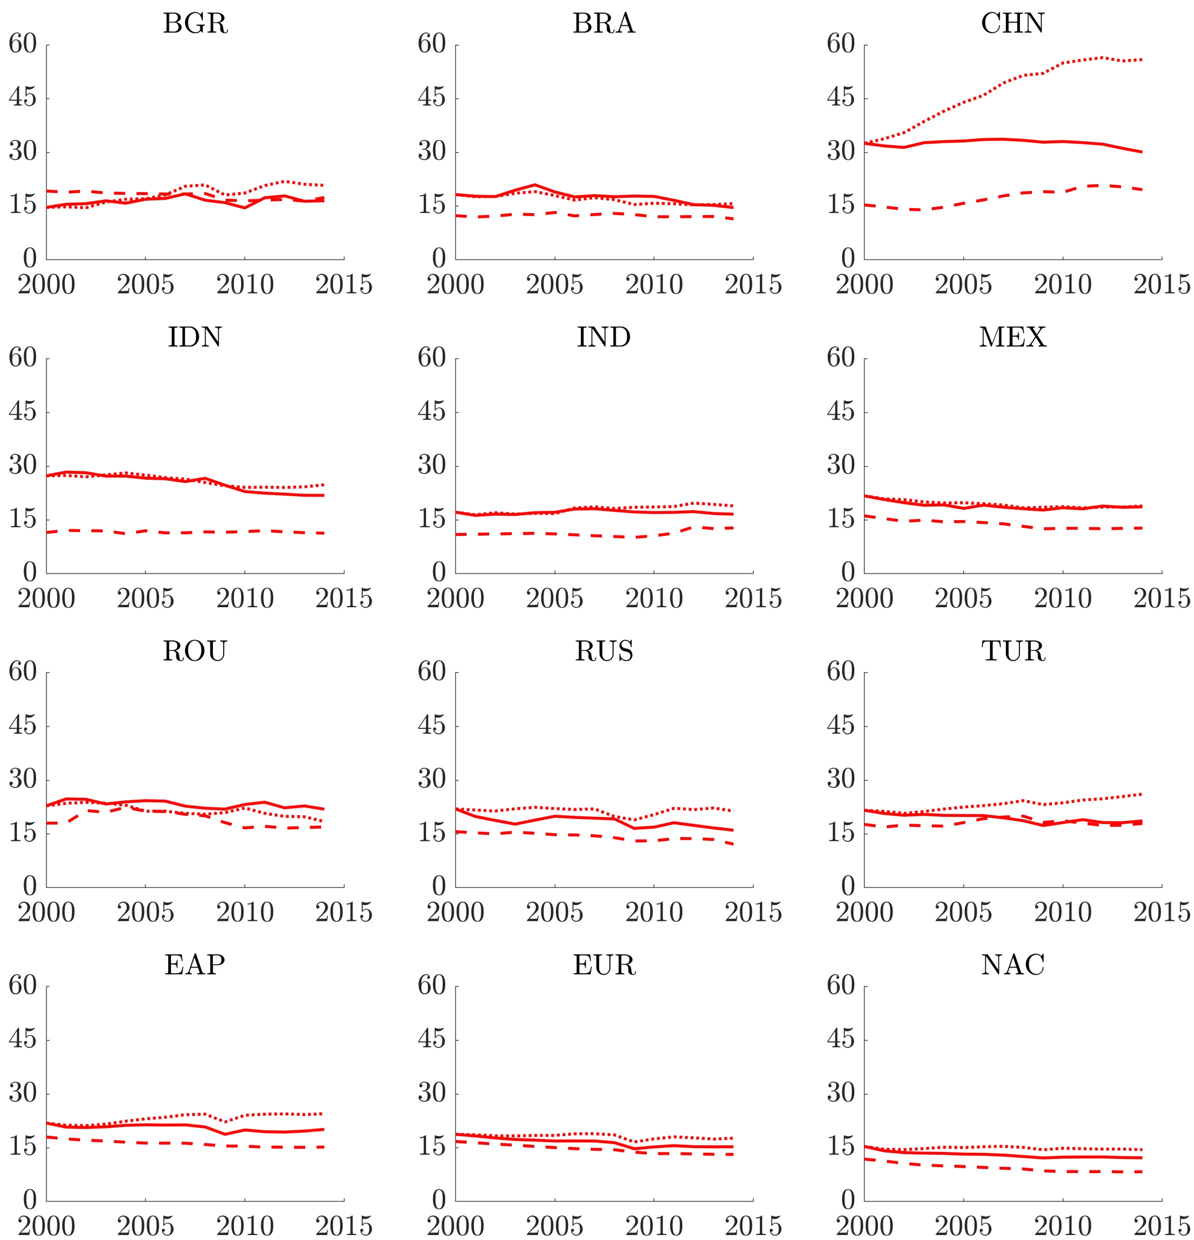

Notes: This figure illustrates three measures of the manufacturing share from 2000 to 2014 for each country/region in the model. The solid, dotted and dashed lines represent the nominal value added share as a percent of GDP, the real value added share as a percent of GDP, and the sectoral employment share, respectively. See Tables C.1 and C.2 in Appendix C.1 for details on country/region codes and sectoral aggregation.
Source: Timmer et al. (2015); Author's calculations.

Figure C.8: Trade Balance (% of GDP)

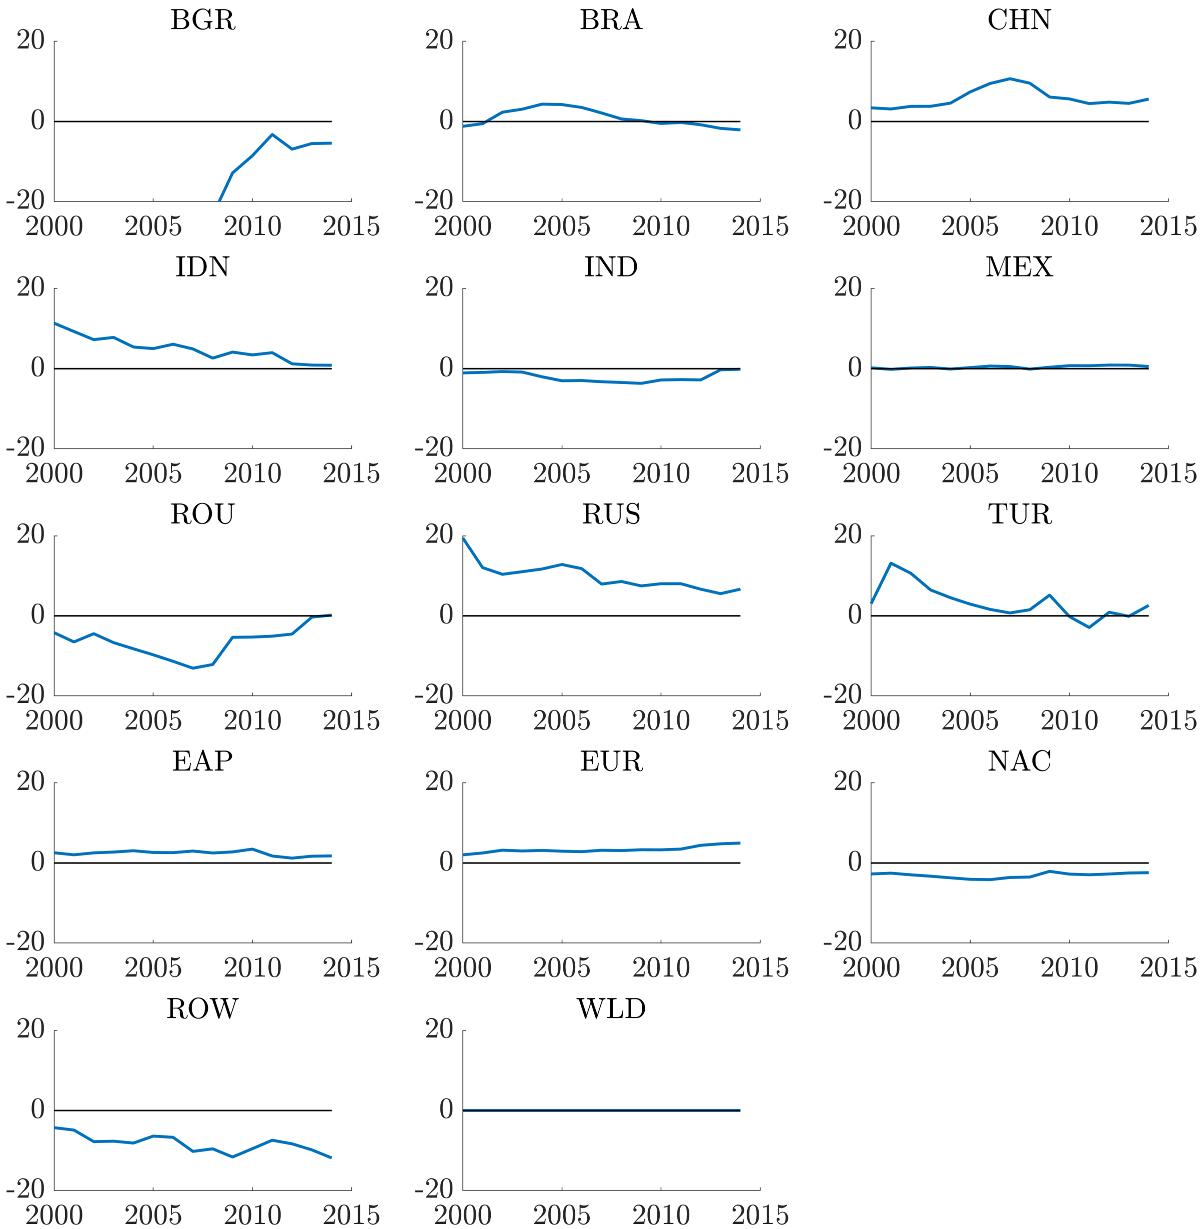

Notes: This figure illustrates trade balance as a percent of GDP from 2000 to 2014 for each country/region in the model. From 2000 to 2008, Bulgaria (BGR) had a trade deficit that ranged from a minimum of 23.4% to a maximum of 31.6% of GDP; these values are omitted from the first panel of the figure due to scaling issues. See Tables C.1 and C.2 in Appendix C.1 for details on country/region codes and sectoral aggregation. Source: Timmer et al. (2015); Author's calculations.

Figure C.9: Trade Balance by Sector (% of GDP)

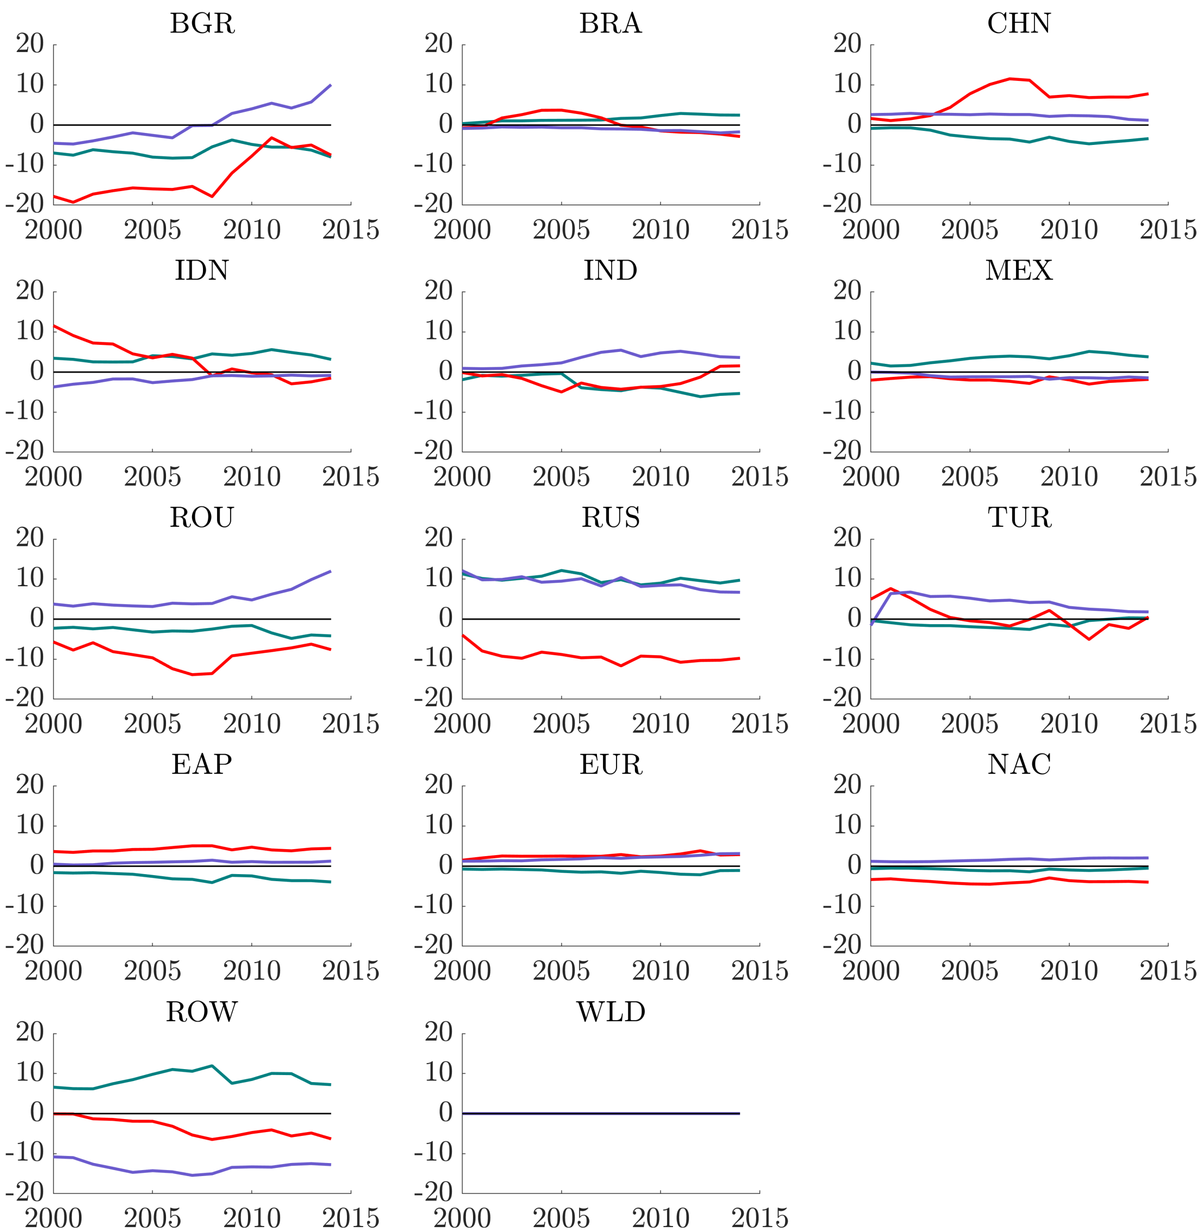

Notes: This figure illustrates the trade balance by sector as a percentage of GDP from 2000 to 2014 for each country/region in the model. The green, red and purple lines represent the trade balance of the Agriculture, Manufacturing and Service sectors, respectively. See Tables C.1 and C.2 in Appendix C.1 for details on country/region codes and sectoral aggregation.
Source: Timmer et al. (2015); Author's calculations.

C.5 Counterfactual Results

The following tables show the changes in the nominal share of value added by sector caused by imposing the various counterfactual exercises. The first six columns of each table are the same. The first three columns show the baseline nominal value added share by sector expressed as a percentage. The next set of three columns shows the baseline trade balance by sector expressed as a percentage of GDP. The final three columns in each table shows the percentage point changes in the nominal value added share of each sector for the various counterfactual exercises. Table C.3 shows the results from shutting down trade in agriculture. Table C.4 shows the results from shutting down trade in manufacturing. Table C.5 shows the results from shutting down trade in services. Table C.6 shows the results from shutting down trade imbalances. Table C.7 shows the results of the interactions effects from imposing all of the previously mentioned counterfactuals simultaneously (that is, the interaction effects from moving to autarky). Table C.8 shows the results from moving to autarky. In all tables, A, M and S are abbreviations for agriculture, manufacturing and services, respectively.

Table C.3: Eliminating Trade in Agriculture - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	0.1	0.8	-0.9
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	0.4	0.3	-0.7
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	-0.5	-0.7	1.2
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	-0.4	0.0	0.3
IND	16.7	16.6	66.7	-5.3	1.6	3.7	-0.1	0.2	-0.1
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	-0.1	-0.1	0.2
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	-0.2	0.3	-0.1
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	-1.5	-0.3	1.8
TUR	9.8	18.6	71.6	0.3	0.5	1.8	-0.3	-0.6	0.9
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	-0.2	-0.2	0.4
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	-0.2	-0.7	0.9
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	0.2	0.6	-0.8
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	1.2	0.7	-1.9
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	0.0	0.0	0.0
WLD	8.9	17.0	74.1	0.0	0.0	0.0	0.0	-0.2	0.2

Table C.4: Eliminating Trade in Manufacturing - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	-0.7	2.4	-1.8
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	-0.8	1.1	-0.3
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	0.6	-2.7	2.1
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	-0.3	0.6	-0.3
IND	16.7	16.6	66.7	-5.3	1.6	3.7	0.5	-0.6	0.1
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	-0.2	0.9	-0.7
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	-0.2	3.3	-3.1
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	-3.0	4.3	-1.4
TUR	9.8	18.6	71.6	0.3	0.5	1.8	0.3	-0.2	0.0
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	-0.4	1.0	-0.6
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	0.5	-2.2	1.7
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	0.1	-1.3	1.2
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	-0.1	1.9	-1.8
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	-0.1	1.9	-1.8
WLD	8.9	17.0	74.1	0.0	0.0	0.0	-0.1	-0.2	0.2

Table C.5: Eliminating Trade in Services - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	2.8	2.3	-5.0
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	-0.3	-0.6	0.9
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	0.4	0.2	-0.6
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	0.0	-0.5	0.4
IND	16.7	16.6	66.7	-5.3	1.6	3.7	1.7	0.6	-2.4
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	-0.3	-0.7	0.9
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	1.8	4.3	-6.1
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	2.8	1.1	-3.9
TUR	9.8	18.6	71.6	0.3	0.5	1.8	0.4	0.6	-1.0
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	1.0	0.8	-1.9
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	0.4	0.3	-0.7
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	0.4	1.3	-1.7
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	0.5	0.8	-1.3
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	-3.6	-2.5	6.1
WLD	8.9	17.0	74.1	0.0	0.0	0.0	0.0	0.2	-0.2

Table C.6: Eliminating Trade Imbalances - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	0.1	0.8	-0.9
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	0.4	0.3	-0.7
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	-0.5	-0.7	1.2
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	-0.4	0.0	0.3
IND	16.7	16.6	66.7	-5.3	1.6	3.7	-0.1	0.2	-0.1
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	-0.1	-0.1	0.2
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	-0.2	0.3	-0.1
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	-1.5	-0.3	1.8
TUR	9.8	18.6	71.6	0.3	0.5	1.8	-0.3	-0.6	0.9
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	-0.3	0.0	0.3
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	-0.2	-0.2	0.4
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	-0.2	-0.7	0.9
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	0.2	0.6	-0.8
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	1.2	0.7	-1.9
WLD	8.9	17.0	74.1	0.0	0.0	0.0	0.0	0.0	0.0

Table C.7: Interaction Effects - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	-0.4	-1.1	1.5
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	0.7	-0.7	0.0
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	-0.7	2.0	-1.3
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	0.5	-0.7	0.3
IND	16.7	16.6	66.7	-5.3	1.6	3.7	-1.6	0.5	1.1
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	0.4	-0.9	0.5
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	-0.7	-3.6	4.3
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	2.1	-3.1	0.9
TUR	9.8	18.6	71.6	0.3	0.5	1.8	-0.3	0.1	0.2
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	0.0	-0.8	0.8
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	-0.7	1.4	-0.7
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	-0.4	-0.1	0.4
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	-0.1	-1.5	1.6
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	2.0	-0.2	-1.8
WLD	8.9	17.0	74.1	0.0	0.0	0.0	0.0	0.0	0.0

Table C.8: Autarky - Change in Value Added Shares by Sector

	Baseline: VA Shares			Baseline: Trade Bal.			Change in VA Shares		
	A	M	S	A	M	S	A	M	S
BGR	7.8	16.5	75.7	-8.0	-7.5	10.1	5.2	1.9	-7.2
BRA	8.8	14.6	76.6	2.5	-2.9	-1.7	-1.4	0.9	0.5
CHN	14.9	30.1	55.0	-3.4	7.8	1.2	1.4	-2.2	0.8
IDN	23.6	21.9	54.5	3.2	-1.5	-0.8	-2.1	1.0	1.1
IND	16.7	16.6	66.7	-5.3	1.6	3.7	4.0	-0.7	-3.2
MEX	10.4	18.7	70.9	3.8	-1.8	-1.4	-2.6	0.9	1.7
ROU	6.6	21.9	71.5	-4.2	-7.6	12.0	2.9	3.3	-6.1
RUS	13.3	16.0	70.7	9.7	-9.8	6.7	-5.2	5.1	0.2
TUR	9.8	18.6	71.6	0.3	0.5	1.8	-0.1	0.1	0.0
LMIC Ave.	12.4	19.4	68.1	-0.2	-2.4	3.5	0.2	1.1	-1.4
EAP	3.2	20.2	76.6	-3.9	4.5	1.3	2.0	-1.9	-0.1
EUR	2.7	15.3	81.9	-1.1	2.9	3.2	0.6	-0.9	0.3
NAC	4.4	12.2	83.4	-0.5	-4.0	2.1	0.8	1.8	-2.5
ROW	22.0	13.6	64.4	7.2	-6.3	-12.8	-4.4	1.2	3.2
WLD	8.9	17.0	74.1	0.0	0.0	0.0	-0.1	0.1	0.0

C.6 Regression on Model Sample

Table C.9: Regression Results: Manufacturing Share of VA (1970-2019)

	Nominal Share		Real Share	
ln population	66.96*	(21.04)	2.33	(18.87)
ln population squared	-1.46**	(0.58)	0.10	(0.52)
ln GDP per capita	63.24*	(8.52)	38.75*	(7.23)
ln GDP per capita squared	-3.56*	(0.49)	-2.02*	(0.41)
1980s	-1.70**	(0.73)	-1.31**	(0.65)
1990s	-8.30*	(0.89)	-2.80*	(0.77)
2000s	-11.35*	(1.08)	-2.83*	(0.94)
2010s	-12.30*	(1.56)	-3.94*	(1.29)
Country/Region F.E.	Yes		Yes	
Countries/Regions	13		13	
Observations	650		650	

Notes: This table shows the results of running the regression specification in (3.1) using the country/region aggregation scheme used in the calibrated model (see Table 158 for details).
Source: Feenstra et al. (2015); United Nations (2021); Author's calculations.