4

BANCO DE PORTUGAL
ECONOMIC STUDIES

4

Banco de Portugal Economic Studies

Volume XI

Please address correspondence to
Banco de Portugal, Economics and Research Department
Av. Almirante Reis 71, 1150-012 Lisboa, Portugal
T +351 213 130 000 | estudos@bportugal.pt

BANCO DE PORTUGAL
EUROSYSTEM

Lisboa, 2025 • www.bportugal.pt

Banco de Portugal Economic Studies | Volume XI – no. 4 | Lisbon 2025 • **Banco de Portugal** Rua do Comércio, 148 | 1100-150 Lisboa • www.bportugal.pt • **Edition** Banco de Portugal • **Design** Communication and Museum Department | Design Unit • **ISSN** (*online*) 2183-5217

Content

Editor's note[1]

Pedro Duarte Neves

October 2025

---

1. This issue of the *Banco de Portugal Economic Studies* includes three studies. The first places the conditions of efficiency and competitiveness of Portuguese firms in the European landscape, through indicators for the use of resources and the outcome of the production process. The second, based on the literature on growth-at-risk, illustrates how financial conditions can contribute to extreme economic downturns. The third study examines the factors that determine free-riding incentives at country level in the greenhouse gas (GHG) emission reduction process.

2. The recent report prepared by Mario Draghi for the European Commission presents several recommendations aimed at reducing the productivity gap between the US and EU economies. This gap is largely the result of an innovation gap, mainly concentrated in the technology sector. At the same time, the possibility of exploiting gains from a deepening of the Single Market has been signalled, in a context where large differences remain between firms in Member States and setbacks in various dimensions of international trade. Comparisons between firms from different countries are therefore essential for conducting informative diagnoses.

3. The study opening this issue of the *Banco de Portugal Economic Studies*, prepared by Amador and Nogueira, analyses the conditions of efficiency and competitiveness of Portuguese firms in the European Union. The comparison uses a set of indicators that characterise the production process – wage, capital ratio, capital per worker, labour productivity and return on assets – for the period 2010-2020. Firms are classified by size class into large, medium and small.

For each indicator, the relative distance is calculated between the median Portuguese firm and the most efficient European firms, identified as those in the highest 75^{th} percentile for the 20 countries considered. The main conclusions of the study are the following, where the indicators are ordered from the most distant to the closest in relation to the most efficient European companies:

E-mail: pneves@bportugal.pt

1. The analyses, opinions and conclusions expressed in this editorial are entirely those of the editor and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.

(i) labour productivity has been disappointing: apart from showing the largest relative difference to the most efficient European firms, this gap has not narrowed in the decade analysed for any of the firm size classes;

(ii) the capital per worker ratio, as well documented in empirical literature, is markedly lower than that of more efficient European firms, a distance which remained, in general, in the decade under review for the three size classes;

(iii) the wage paid by the median Portuguese firm has become distant from the wage paid in the most efficient European firms, across all firm size classes and in most economic sectors;

(iv) return on assets has moved somewhat closer to the most efficient standards at European level;

(v) the capital ratio is the variable closest to the most efficient European standards, and this difference has narrowed over the period considered for the three firm size classes.

The study presents some interesting conclusions:

(vi) during the period considered, large Portuguese firms were generally more efficient in relative terms than small and medium-sized, as they are closer to the most efficient European firms in their size class;

(vii) for the indicators considered, Portuguese firms are close to, although generally below, the median of the 20 countries considered, with a position prevailing in the lower quartile when considering those countries which, like Portugal, participated in the initial stage of the euro.

4. There is extensive evidence in the economic literature that economic recessions accompanied by financial crises tend to be more severe – in terms both of duration and contraction in economic activity – than economic recessions without financial distress. In the aftermath of the Great Financial Crisis, the high costs of financial crises justified the development of macroprudential policy as an essential element in preserving financial stability.

The empirical distribution of observed rates of change in output shows, for most advanced economies, a negative asymmetry with heavy tails (especially the left-hand tail). Literature also highlights the existence of non-linearities in the relationship between the distribution of future output growth and financial indicators.

The growth-at-risk (GaR) analytical tool seeks to integrate this empirical evidence, as it essentially relates the prevailing financial conditions in the economy to the distribution of future GDP growth. Literature matches the lowest percentiles of this distribution (generally the 5^{th} percentile) to GaR, focusing on the most adverse economic situations.

The GaR methodology presents several features that justify their use in conducting macroprudential policy:

(i) the GaR methodology makes it possible to assess the entire distribution of future growth of the economy, representing a considerable advantage over the exclusive use of the central location measures of that distribution (such as forecasting average growth or modal growth) typical of macroeconomic projections. In addition, the focus of the GaR methodology is on the left-hand tail of the distribution of future output growth, which precisely corresponds to the most severe downturns;

(ii) Considering the distribution of future GDP growth in the medium term (which goes beyond next year) allows for the capturing of key aspects for financial stability. For example, more favourable financial conditions in the short term – translated into lower risk premia and higher credit growth – will tend to result in higher output growth during that term. However, from a medium-term perspective, if the economy is accumulating structural imbalances that may be corrected more abruptly at a later stage of the economic cycle, the conditions for a crisis with predominantly financial contours may be created;

(iii) The GaR approach also makes it possible to assess the impact of a build-up of systemic risk on future output behaviour, thus being a key analytical element for macroprudential policy. In particular, it offers the possibility to assess the trade-off that exists between output growth and the GaR in conducting macroprudential policy: more (less) restrictive macroprudential policy will tend to have a negative (positive) effect on output growth in the short term, but it will also tend to contribute to lower (higher) severity if a particularly adverse situation occurs (indicated, for example, by the 5th percentile of the distribution of future output growth).

5. The second study of this issue of the *Banco de Portugal Economic Studies*, prepared by Buratta and Lima, applies the GaR methodology to the Portuguese economy. Through a New Keynesian model incorporating a financial vulnerability channel, which allows non-linearities between financial conditions and future output growth to be captured, the authors can show that in relation to Portugal for the period 1994-2019:

(i) the probability of extreme achievements of the output gap (associated with the most severe economic downturns and thus on the left-hand tail of the distribution) responds to worsening financial conditions;

(ii) the volatility of distribution of the change in the output gap – as measured by the difference between the volatility of the 5th percentile and that of the 95th percentile – is not constant and is negatively correlated with the median;

(iii) excessive risk-taking that characterises periods of low economic volatility tends to result in a worsening of the risk of greater future instability.

6. On 29 November 2007, the British economist Nicholas Stern delivered a public lecture at the Royal Economic Society, which was a key milestone in applying economic theory to climate change.[2] It is therefore justified to transcribe some essential messages from this public lecture:

"The problem of climate change involves a fundamental failure of markets: those who damage others by emitting greenhouse gases generally do not pay. Climate change is a result of the greatest market failure the world has seen. The evidence on the seriousness of the risks from inaction or delayed action is now overwhelming. (...) The problem is global and the response must be a collaboration on a global scale. (...) targets and trading must be at the heart of a global agreement to reduce greenhouse gas emissions".

Around three years later, in January 2011, the World Economic Forum recognised the importance of climate risks for the first time in the Global Risk Report 2011. The Paris Agreement, signed on 12 December 2015, represented a crucial event in recognising the need for a planetary approach to tackle climate change.[3]

After almost 10 years, the latest geopolitical developments, including the US withdrawal from the Paris Agreement, raise concerns about the role of international cooperation in global climate action. For that reason, assessing what the results of non-cooperative solutions might be, and when there are free-riding behaviours is also justified.

7. GHG emissions – the main cause of global warming and climate change – are a negative externality, as private emission costs do not correspond to social costs. The answer of economic theory to this problem is clear: emitters must pay a price for carbon emissions, with social costs internalised in the decision process. Government intervention can take two forms: taxation (carbon taxes) and the creation of a market for emission allowances.[4]

2. In 2007, Cambridge University Press also published the book by Nicholas Stern, "*The Economics of Climate Change: The Stern Review*", which addresses the following topics: impact of climate change on growth and development, policy responses for mitigation and adaptation, and the need for coordinated international action.

3. The Paris Agreement, as is well known, is a historic multilateral commitment to limit the increase in the global average temperature well below 2ºC in relation to pre-industrial levels and to pursue efforts to limit this increase to 1.5ºC.

4. These two possibilities are well studied by economic science, with the seminal works from UK economists Arthur Pigou and Ronald Coase.

Climate is a planetary public good. Thus, the only socially optimal approach to this market failure would be a concerted global response. The impossibility of achieving uniform global carbon taxation or developing a planetary market for emission allowances inevitably creates free-riding incentives: setting national emission reduction targets and thus emission levels above the overall optimal level. It is therefore essential to understand these free-riding behaviours – more pronounced, as already mentioned, in the current context of geopolitical developments that draw effectiveness from international cooperation – and their practical consequences.

The final study of this issue of the *Banco de Portugal Economic Studies*, prepared by Adão, offers several theoretical and applied contributions to a better understanding of this reality, with reference to the introduction of carbon taxes. The main theoretical findings are as follows:

(i) the overall optimal solution is a cooperative balance where all countries apply an equal Pigouvian rate – corresponding to the environmental damage from emissions – that fully internalises the overall externality caused by emissions;

(ii) in the case of a non-cooperative balance, each country determines its carbon rate, which tends to be lower than the overall optimal rate (the Pigouvian rate);

(iii) the size of the country affects the free-riding incentives: all other things being equal, smaller countries have fewer incentives to cut emissions than larger countries;

(iv) The severity of climate damage also affects the free-riding incentives: all other things being equal, countries more affected by climate change tend to opt for higher carbon rates than less affected countries.

Adão also presents a case study, which is not particularly plausible but still offers potentially useful findings, and which seeks to assess the hypothetical situation of China being the only country with free-riding behaviour, i.e. the only country not applying the optimal overall Pigouvian rate. The main findings are as follows:

(v) as it is a large economy – see finding (iii) above –, China would have incentives to apply a carbon taxation closer to but below the optimal overall Pigouvian rate;

(vi) for that reason, this free-riding behaviour would cause overall losses that are unavoidable as a deviation from the planetary optimal, which would be partially mitigated by the fact that, due to its size, China sets a carbon rate relatively close to the Pigouvian optimal.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. XI, No. 4

Non-technical summary

October 2025

Inputs and outcomes of portuguese firms: A micro-aggregated EU comparison

João Amador and Guida Nogueira

Recent decades have been characterized by the increasing availability of firm-level data and subsequent progress regarding comparisons across firms. Such empirical research has produced useful insights on productivity drivers, which are necessary to address market failures and inform public policies and investments.

Although the literature that looks into firms' characteristics in single country has greatly expanded, the comparison of firm-level dimensions across countries is quite rare. Our article aims to contribute in this firm-level cross-country dimension by measuring the relative distance of the median portuguese firm to the best performing ones in the set of 20 European Union (EU) countries, and covering the period 2010-2020. Such relative distance is computed for a set of five indicators, classified as input (average wages, equity ratio and capital per worker) and outcome variables (labour productivity, measured as the gross value added per worker, and return on assets – ROA).

As for the input variables, we observe that the median portuguese firms pay relatively low wages when compared with the firms in the highest 75th percentile of the EU distribution, and their position is intermediate in the ranking of the twenty EU countries considered. This distance has visibly increased in the period studied, and it is present in all firm sizes and sectors. As for the equity ratio, the distance of the portuguese median firms versus the best performers is in an intermediate position and there was a reduction in this gap in the period 2010-2020. Moreover, there is a sizable distance to the firms in the top of the distribution in terms of the capital per worker ratio, which has not reduced in the last decade. Nevertheless, the distance of the median firms to the best performers in the EU is high in many countries and the median portuguese firms are in an intermediate position in the ranking.

As for the outcome variables, there has been some increase in the relative distance to the firms with the best performance in the EU in terms of labour productivity, with small firms comparing worse than their medium and large counterparts. Finally, in terms of ROA, the results are more favourable for the largest firms. This group shows a smaller distance to the EU's best performers and holds a better position in the overall ranking when compared to firms in other size classes.

Overall, the median portuguese firm in each size class and sector exhibits a significant performance gap relative to the best performing firms in the EU in the corresponding category. This is neither a surprising result nor exclusive to Portugal, because the distributions for the variables considered are typically strongly right skewed. Nevertheless, this distance is acute in the capital-labour ratio, average wages and labour productivity. The path of the relative indicators in the period 2010-2020 shows some convergence toward the best performing EU firms in the equity ratio and widening gaps in terms of average wages and labour productivity. The somewhat disappointing path of labour productivity is consistent with results obtained using a similar methodology but based on national accounts aggregate data.

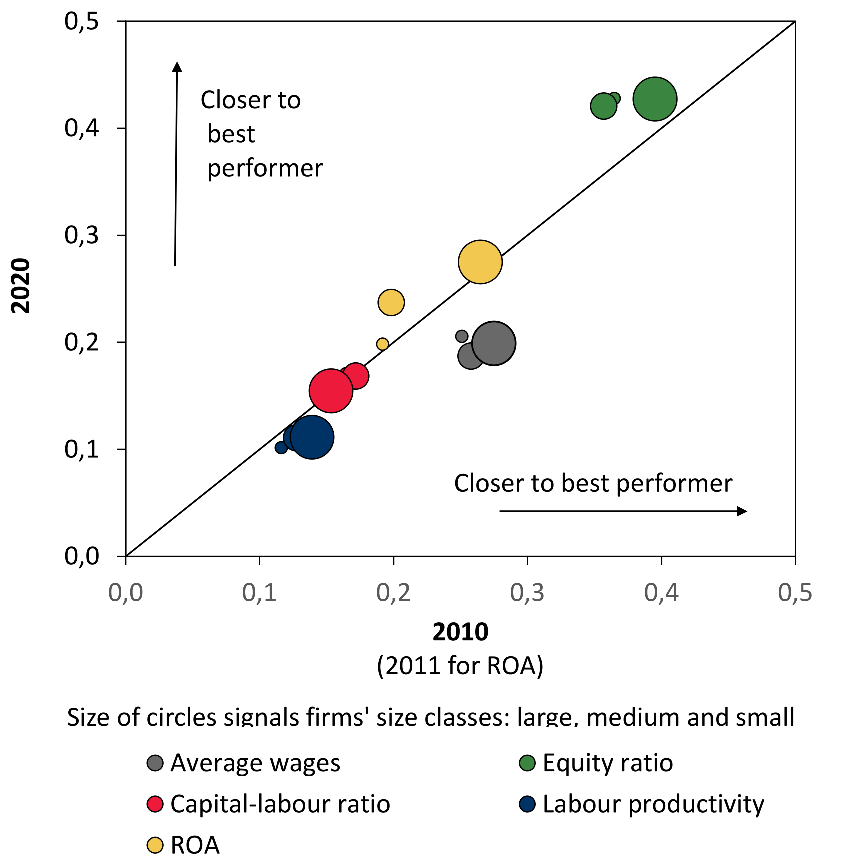

FIGURE 1: Relative distance of the median portuguese firms to the best performers (75th percentile) in the EU

Note: The metric used for comparison is the relative distance of the median firm to the firms with the best performance in the sample, considered as those in the highest 75th percentile among the 20 EU countries considered. The distance is standardized for the difference between the highest 75th percentile and the lowest 25th percentile in the EU distributions, thus ranging in the zero-one interval. Since the distributions for the variables considered are strongly right skewed, the metric takes lower values in this interval. Circles above (below) the diagonal mean an improvement (deterioration) from 2010 to 2020.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. XI, No. 4

Inputs and outcomes of portuguese firms: A micro-aggregated EU comparison

João Amador
Banco de Portugal
Nova SBE

Guida Nogueira
Portuguese Ministry of Economy

October 2025

Abstract
The analysis of firms' inputs and outcomes is a well-established empirical exercise, which has been made possible by the increasing availability of firm-level data. However, the comparison of firm-level dimensions across countries is quite rare. This article aims to contribute in this dimension, focusing on Portugal within a set of 20 European Union countries and covering the period 2010-2020. Data constraints impose that the analysis takes only firms with at least 20 employees, which represent the large majority of gross value-added. The metric used is the relative distance of the median firm to the firms with the best performance in the sample, considered as those in the highest 75th percentile among the EU countries considered. The distance is standardized for the difference between the highest 75th percentile and the lowest 25th percentile in the EU distributions. As for input variables, we study average wages, equity ratio, and capital per worker. In what concerns outcome variables, we focus on labour productivity, measured as gross value added per worker, and return on assets. We observe that the median portuguese firm in different sectors and size classes is quite distant from the best performing firms in the EU. This is neither a surprising result nor exclusive to Portugal, because the distributions of the variables considered are typically strongly right skewed. However, the relative performance of the portuguese median firm is somewhat disappointing because its relative distance to the best performers has increased along time in several dimensions, notably in labour productivity. (JEL: O47, O52)

"We will have to learn more on productivity to drive growth."
Mario Draghi, 2024

Acknowledgements: The authors are thankful to Joana Almodovar, Nuno Alves, António Antunes, Paulo Inácio e Pedro Duarte Neves for their useful comments and suggestions. The opinions expressed are those of the authors and do not necessarily coincide with those of Banco de Portugal, the Eurosystem or the Portuguese Ministry of Economy.

E-mail: jamador@bportugal.pt; guida.nogueira@gee.gov.pt

1. Introduction

In the current economic and geopolitical landscape, the European Union (EU) faces the ever greater challenge of improving the performance of its firms to address the persistent and growing aggregate productivity gap versus global peers. This concern has been highlighted in several high-profile policy reports (Draghi 2024; Letta 2024) and translated into an EU road-map to "regain competitiveness and secure sustainable prosperity" (European Commission 2025). These reports highlight that the core of a competitiveness agenda should be higher productivity growth and acknowledge the strong heterogeneity between EU firms in terms of inputs and performance. Beyond the policy approach, research about European productivity has been progressing (Bergeaud 2024; Adilbish *et al.* 2025). As for Portugal, this debate has also been very active, motivated by the comparatively low levels of output per worker in the EU context. Examples of such studies for the portuguese economy are Arnold and Barbosa (2015), CPP (2019) and Banco de Portugal (2019).

The meaning of *competitiveness* and its causal interplay with *productivity* has always been a source of discussion in the literature and policy debates. It is widely acknowledged that competitiveness is interpreted as having the conditions for a sustained good performance versus other players, but whether it relates most to countries as a whole or firms is debatable. In market economies, firms are born, survive, and thrive by competing successfully in the markets (Connor 2003). Therefore, the performance of firms is often discussed through a lens of competition, with the latter force separating winners from losers. In this vein, (Krugman 1994, 1996) claims that "*competitiveness is a meaningless word when applied to national economies*", and it should only apply to individual firms. However, this perspective ignores the interconnectedness between the micro, meso and macro levels of economic activity (Peneder 2017). Countries compete economically, but in a different way from firms (Dunn 1994). For the most part, countries' competition is indirect, by providing favorable business conditions that create and shape firms' competitive advantages (Peneder 2017).

In Amador *et al.* (2022) we take the view that competitiveness corresponds to a set of aggregate conditions that are likely to sustain a good economic performance relative to the benchmark countries, measured as the relative output per worker in the overall economy. We believe that having a wide set of comparable firm-level variables for a relatively long time span would make it possible to develop a firm-level competitiveness analysis, as presented in Martins *et al.* (2022) for five EU countries, using the iBACH database, developed under the aegis of the European Committee of Central Balance Sheet Data Offices.

Competitiveness can thus be conceptualized and measured along different layers of economic activity to capture specific dimensions of performance, but it will always be a relative concept that assesses performance versus a reference group. Moving the focus from a macro-level approach of competitiveness to a firm-level approach calls for a micro-founded cross-country dataset, as discussed in Bartelsman *et al.* (2023). Even though such international benchmarking remains largely limited due to data confidentiality, different compilation methodologies, and different time horizons, there is progress in terms of coverage, representativeness and cross-country comparability of micro-aggregated data. The Competitiveness Research Network (CompNet) dataset is a good example of such progress.

A major hurdle when the objective is to study competitiveness at the firm level is the very low number of firm-level variables that are comparable across countries and their short time span. The list of variables that potentially determines firms' competitiveness is long and includes information about the quality of inputs and not just on their quantity. For example, the skills of the workforce, the characteristics of the management team, the types of investment made, including R&D and patents, the market strength of the brand, the environmental and quality certification of products and so on. Accessing this type of granular data is very hard, but to have it at a comparable level across countries is presently virtually impossible, and still distant in time. Therefore, researchers have to accept the fact that only a very limited list of variables is available.

A related aspect concerns the intrinsic noisy nature of firm-level data. The distribution of firms' size and productivity is typically Pareto-like, with a large number of micro firms, sometimes with negative numbers for GVA, and a small number of extremely high productivity firms. Therefore, if the objective is to assess distance to best performers and standardize by the distance between the best and worst performers, it is almost certain that we are comparing outliers. Therefore, it is necessary to define best (worst) performers as firms in intermediate upper (lower) percentiles of the distributions. This choice obviously involves a degree of discretion and is made harder if the cross country databases are not fully comparable in terms of their samples.

Shifting from EU-wide averages to country-specific distributions, from aggregate to micro-aggregated dynamics, and from absolute distances to relative metrics, offers an insightful perspective on the determinants of productivity gaps. In this vein, the objective of this article is to contribute to a deeper understanding of the position of portuguese firms, relative to their peers in the EU, focusing on a set of three input variables – average wages, equity ratio, and capital per worker – as well as on two outcome variables – labour productivity, measured as GVA per worker, and return on assets (ROA).

We observe that the median portuguese firms pay relatively low wages when compared with the firms in the highest 75th percentile of the EU distribution, and their position is intermediate in the ranking of the twenty EU countries considered. This distance has increased in the last decade and it is present in all firm sizes and sectors. As for the equity ratio, the distance of the portuguese median firms versus the best performers is in an intermediate position and there was a reduction in this gap in the last decade. Moreover, there is a sizable distance to the firms in the top of the distribution in terms of the capital per worker ratio, which has not reduced in the last decade. Nevertheless, the distance to the best performers is high in many EU countries and the median portuguese firms are in an intermediate position in the ranking. As for the outcome variables, there has been some increase in the relative distance to the firms with the best performance in terms of labour productivity, and the position in the

EU ranking is lower for small firms. Finally, as for ROA, the distance to the best EU performers and the position in the ranking is somewhat more favourable among the largest portuguese firms.

The article is organized as follows. In the next section we briefly present the data sources, the input and outcome variables, and discuss their shortcomings. Section 3 presents the metric used to set the distance of the median firm to the best performers and the percentiles that define them. Section 4 presents the results and it is organized along two blocks. Firstly, we discuss the input variables, comparing along firm size class, time and main sectors of activity, while comparing across EU countries. Secondly, we focus on the two outcome variables and replicate the structure used for the analysis of inputs. Next, in section 5 we correlate input and outcome variables and compare the micro-aggregate result with an aggregate approach. Section 6 presents some concluding remarks.

2. Data

The article uses the latest vintages of the Competitiveness Research Network (CompNet) dataset, an unbalanced panel that contains detailed micro-aggregated information in a harmonized, cross-country setup. The dataset is compiled in collaboration with a network of national data providers, mostly National Statistical Institutes or National Central Banks, through a distributed micro-data approach that applies common standardized statistical routines on firm-level data, representative of the entire population of non-financial corporations in each participating country. The resulting dataset provides a set of micro-founded variables classified by firm size, sector, country and year. Variables are clustered in six main topics: labour, finance, competition, trade, productivity, and other. For each variable, a comprehensive set of statistical moments is available, namely the 1st, 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles, alongside the mean, standard deviation, skewness and kurtosis. Therefore, we do not observe the percentiles in the distribution of EU firms, but only the distributions in each country, thus explaining our pick of the largest 75th percentile and lowest 25th percentile as benchmarks. The initial study using this dataset is CompNet (2014).

Although the time span and indicator availability vary across countries and time windows, the 10th and 9th vintages of the dataset together alow us to build a panel with 20 EU countries, spanning the period between 2010 and 2020. Table A.1 in Appendix A lists the countries considered and proxies the coverage of the CompNet database relatively to the Eurostat Structural Business Statistics (SBS) in terms of number of firms and GVA. In terms of sectoral coverage, the dataset includes nine macro-sectors, corresponding to NACE Rev.2 sections, consistent with the European System of Accounts categories. Table A.2 in Appendix A contains a detailed definition of the macro-sector and industry-level aggregation. Furthermore, the dataset covers non-financial corporations with at least one employee following the same size-class definitions of Eurostat SBS. Table A.3 in Appendix A contains the size-class definitions used in the article. Additionally, because firm reporting requirements vary across participating countries, often associated with different reporting thresholds, the database presents two samples: one that includes all firms in the target population with at least one employee and another that includes only firms with at least 20 employees.

To set the stage for an effective benchmarking exercise, three key decisions underpin our analysis. Firstly, to improve cross-country comparability and ensure a balanced panel over time, the analysis is limited to a subset of 20 EU countries with the 2010-2020 time frame. Secondly, the real estate macro-sector is also excluded from the analysis due to its distinct structural characteristics. Finally, to mitigate the impact of reporting differences and ensure a more homogeneous sample for cross-country comparisons, micro-firms are excluded from the analysis. This decision is important because it means dropping a very large number of micro firms, whose different performance across countries could bias our conclusions. Nevertheless, although firms in our sample are just a small share of the total, they represent the large majority of the GVA in the economy (see Table A.1 in Appendix A).

3. Methodology

The methodological approach adopted in this article is based on the distance of a set of firms' input and performance indicators relative to the best performers within the EU sample. This measurement is computed in a comparable and standardized way by using a suitable metric, and with a focus on Portugal. The sequence of the analysis for each input and performance variable is: i) analyze the evolution of the relative distance of the portuguese median firm to those with the best results in the EU in the period 2010-2020; ii) compare the relative distance for different sectors of activity in 2010 and 2020; iii) rank the countries in the sample according to the relative distance of their median firms in 2020; iv) identify the position of the median portuguese firm in the histogram of the cross-country distribution. The histogram is complementary to the ranking of the countries because it highlights the position of the median portuguese firm in the distribution. In fact, a low position in the ranking can still correspond to a good performance if all other countries perform well, and the opposite is also possible.

A meaningful comparison of firms' inputs and outcome variables requires controlling for their basic characteristics, notably size and sector of activity. Therefore, we separately report results for small, medium, and large firms. Moreover, calculations are made at the basic sectoral level and aggregated with the weights of each sector on the country's total gross value added (GVA) in the average of the period 2010-2020.

3.1. Metric

Comparing across different variables and along time requires a common metric and the elimination of scale effects. For this purpose we chose the so-called "min-max" normalization method. For each country p, this method expresses all indicators (I_ip) in a common range, from 0 (furthest from the best performance) to 1 (closest to the best performance). For each basic variable x_ip, the transformation "min-max" takes the distance of the median firm relative to the best performing firm in year t, normalizing this distance by the amplitude between the best and worst performers in that year.

Although this transformation enables comparisons over time, it is quite sensitive to the existence of extreme observations. In order to overcome this problem we do not consider the maximum and minimum of each variable in each year, but instead its largest 75th ($maxP75(x_i^t)$) and lowest 25th ($minP25(x_i^t)$) percentiles among the set of EU countries in the sample. Equation 1 defines our metric and Figure 1 presents a stylized visualization with three countries.

$$I_{ip}^t = 1 - \frac{maxP75(x_i^t) - P50x_{ip}^t}{maxP75(x_i^t) - minP25(x_i^t)} \tag{1}$$

At first glance, the choice of the largest 75th and the lowest 25th percentiles in the distribution of each variable as a way of setting the best or worst performers in the EU may raise concerns. However, as explained in the previous section, since the underlying firm-level data often presents a Pareto-like distribution, the difference between top and low percentiles, and their difference to the median is very large. Therefore, choosing very high or very low percentiles as benchmarks would overstate distances to best performers and make it harder to interpret results.

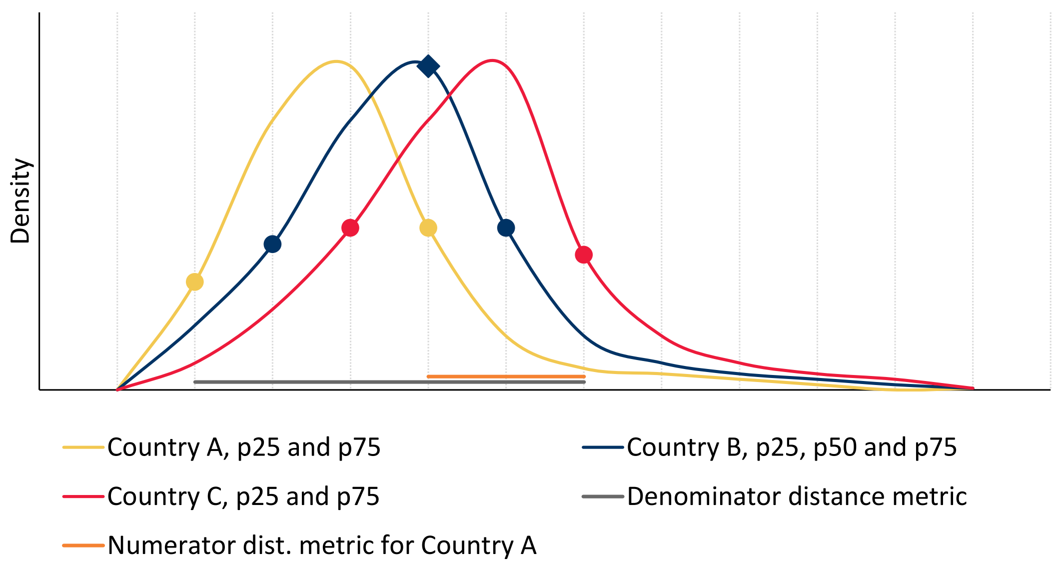

FIGURE 1: Stylized illustration of distance metric

3.2. *Selected variables*

In this subsection we briefly discuss the relevance and interpretation of the input and performance variables analyzed: average wages, equity ratio, capital per worker, GVA per worker and ROA.

3.2.1. *Input variables*

Average wages are defined in the database as total nominal labour costs divided by the number of employees (headcounts). Total labour costs comprise both direct (wages and salaries) and indirect (benefits, social security contributions, recruitment, training) costs. This ratio is an imperfect indicator of firms' underlying cost competitiveness. This input price feeds in the total costs of the firm and affects the price of the final product. However, wages are strongly determined by qualifications, skills and other characteristics of the workers. Unfortunately, we cannot study these dimensions due to the lack of cross-country comparable data. The option of considering the unit labour cost (ULC) indicator as an input variable cost has the disadvantage of embodying the GVA per worker in the calculation, thus mixing input and outcome variables. Although wages are a function of (marginal) labour productivity, it is important to examine them separately.

Equity ratio is defined as the ratio between equity (total shareholders' funds including shares issued, retained earnings, additional paid-in capital, reserves and non-controlling interest) and total assets (sum of current and fixed assets). It is a classic financial indicator that consists of the proportion of firms' assets that are financed by shareholders' equity rather than debt. A higher equity ratio signals greater financial stability and thus a lower risk of bankruptcy. The firm's financial situation is not an input per se. Nevertheless, it conditions the firms' decisions and, ultimately, its performance. There are few studies that use micro data to shed light on the relationship between finance and firms' outcomes. Some studies show that there is a positive relationship between debt and productivity growth, but several new macro-econometric studies have shown that there is a threshold of financial development above which debt negatively impacts growth (Gomis and Khatiwada 2017).

Capital per worker is defined as the value of capital stock (sum of tangible and intangible assets) in real terms (using the implicit price deflator for fixed capital consumption, at NACE Rev.2 two-digit sector level from Eurostat) divided by the number of employees (headcounts). It is a classic ratio, used to assess the position of a firm in a Cobb-Douglas production function, and thus intrinsically related with he GVA per worker (Solow 1957). The capital stock is not homogeneous, and its composition is intertwined with the skills of the workers and the sector of activity. It would be desirable to have a separation between tangible and intangible capital in the firms. However, this separation is not available in the database for all countries considered and in the period considered. Over and above the availability of data lie the difficulties in measuring the capital stock, notably the fact that book value may differ from actual stock of capital actually in operation.

3.2.2. *Outcome variables*

Labour productivity is defined as real GVA (using a sectoral deflator) divided by the number of employees (headcounts). This ratio can be designated as output per worker or *apparent* labour productivity. It is important to highlight that this measure is not adjusted for the human capital of the worker and it is a quite different concept from the total factor productivity (TFP). This latter concept captures the impact of all factors beyond the quantity of inputs, i.e., it does not include the labour productivity.

Return on assets (ROA) is defined as the ratio between operating profit/loss (EBIT) and average total assets over a two-years ($t-1$ and t). It is a classic financial ratio that indicates how profitable a firm is relative to its total assets and is a proxy for the efficiency of the firm. However, there are some caveats to consider. Firstly, the ROA is based on accounting figures, thus different depreciation methods, asset valuations, or leased assets (on or off balance sheet) can skew the ratio. Relatedly, older firms with depreciated assets may have artificially high ROA. Moreover, firms with intangible assets not captured on the balance sheet may appear to underperform. Finally, this ratio is very affected by cyclical developments, which are not necessarily synchronized across sectors and countries. Therefore, it is preferable to assess the path of the ROA across countries and sectors along several years.

As previously referred, there are other indicators that would be useful for a comprehensive analysis of firms' competitive conditions. The list is long and, beyond assessing the quantity and the quality of capital and labour inputs, it would include information on management practices, accessibility to infrastructure and energy, the existence of patents, engagement in international trade and the presence of foreign capital. A summary description of the six variables considered in our study is presented in Table C.1 in Appendix C.

Variable	Description
Input variables	
Average wages	Total labor cost / number of employees
Equity ratio	Equity / total assets
Capital intensity	Real capital stock / number of employees
Outcome variables	
Labour productivity	GVA real / number of employees
Return on assets	EBIT $/(Assets_{t-1} + Assets_t) * 0.5$

TABLE 1. Selected input and outcome variables

4. Results

This section assesses the distance of median portuguese firms to the highest 75th percentile in the distributions of EU firms. We present the results graphically and divide the analysis in five similarly structured blocks - three input variables and two outcome variables. A summary table with the numerical results underlying the graphs is presented in Appendix B and the list of countries-firms that determine the thresholds is provided in Appendix C.

4.1. Average wages

Panel A) of Figure 2 plots the path of our metric when applied to average wages in portuguese firms. The relative distance of median average wages in portuguese firms relatively to that of the highest 75th percentile of EU firms is high and has widened in the period 2010-2020. The distance stood slightly below one third in 2010 and became close to one fifth in 2020. This downward path was common to firms in different size classes, but the indicator stabilized in the latest years for small firms, whose level surpassed the one for medium and large firms. This means that, although very distant from the highest 75th percentile of EU average wages in their size class, portuguese small firms are not more distant from the best performers than their medium and large portuguese counterparts. As mentioned above, this downward path can be read differently. On the one hand, if we take the perspective that in equilibrium wages match labour productivity, this signals an underlying lower performance of the labour input in portuguese firms in the decade 2010-2020. On the other hand, if we adopt the perspective that relatively lower wages provide a competitiveness edge to firms, notably in terms of FDI attraction, and signal labour market flexibility, these developments could be taken more positively.

Panel B) of Figure 2 compares the level of our metric across portuguese sectors in 2020, also comparing the three firm size classes. Differences across sectors are relevant and there is no clear pattern when we compare firm size classes in the same sector. Overall, the sector that is more distant from the highest 75th percentile is Manufacturing, but large firms in Trade & repair and small firms in Construction are also far from the best performers in the EU distribution. Sectors where the average wages of the median firm are closer to the highest 75th percentile in the EU are Transports and Accommodation, but less so for large firms in the latter sector.

The second part of Figure 2 presents a sequence of three panels, one for each firm size class, where EU countries are ranked according to the level of the indicator in 2020, while also signaling the level in 2010. The figures correspond to the weighted average of sectoral levels, using GVA weights. Portuguese firms are placed in the lower half of the ranking for all size classes, and very distant from those with the average wages closer to the 75th EU percentile. These are firms from Sweden, Belgium and Denmark, with relative distances between 0.7 and 0.8.

The fact that the levels of the indicator change markedly along the ranking leads us to present the same information under the form of an histogram, where we signal the bin where the median portuguese firms are and the vertical axis expresses the number of countries in each of them. The three bottom histograms of Figure 2 signal distributions far from uniform and with a mode in the second lowest bin. The position of portuguese medium and large firms is in the mode of the corresponding histograms. As for the small firms, the distribution has comparatively more density on the top bins, and the median portuguese firm stands in the third bin, immediately after the mode.

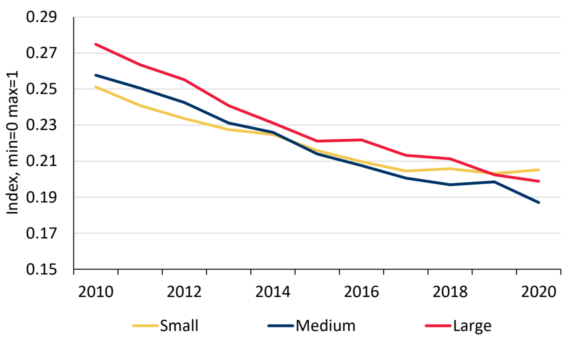

(A) 2010-2020

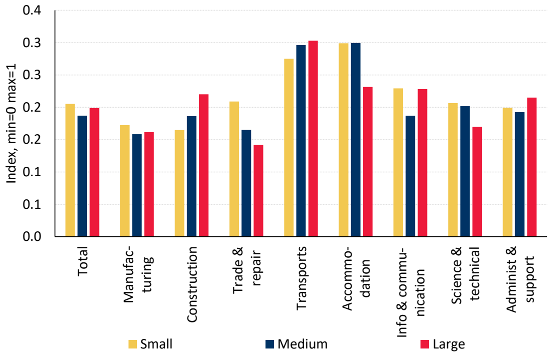

(B) Sectors, 2020

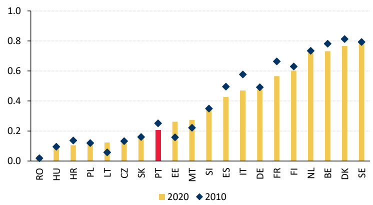

(C) Small firms

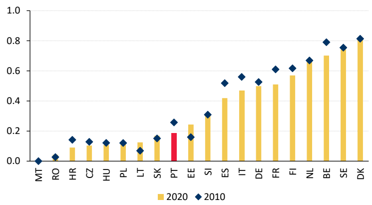

(D) Medium firms

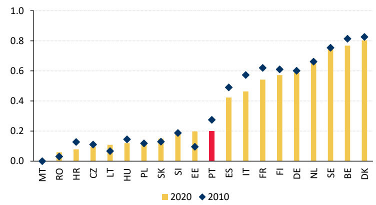

(E) Large firms

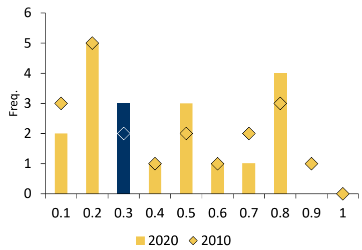

(F) Small firms

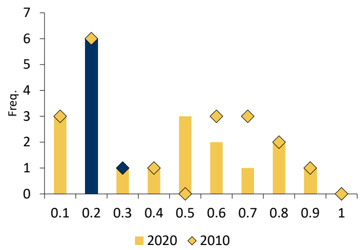

(G) Medium firms

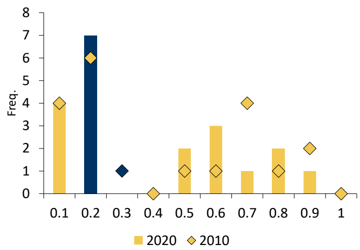

(H) Large firms

FIGURE 2: Average wage: Time, sector, rank across countries and distribution

Note: Blue bars in the histograms identify the bin where the median portuguese firm is.

4.2. *Equity ratio*

As previously said in subsection 3.2, the equity ratio is an indicator of firm's solvency but also signals the development of financial markets and the availability of external financing to the firm. Panel A) of Figure 3 shows a relative increase of the equity ratio of portuguese firms versus those in the highest 75th percentile of the EU distribution, from 2010 up to 2020.

This path is consistent with strong reduction of indebtedness of firms after the portuguese Economic and Financial Assistance Program negotiated with the international institutions (European Commission, ECB and IMF) in the context of the euro area sovereign debt crisis. Both the levels and the path of the indicator are similar for different firm-size classes. Nevertheless, the relative distance to the 75th percentile of the distribution is larger than 0.5 in 2020. As previously stated, since we set as

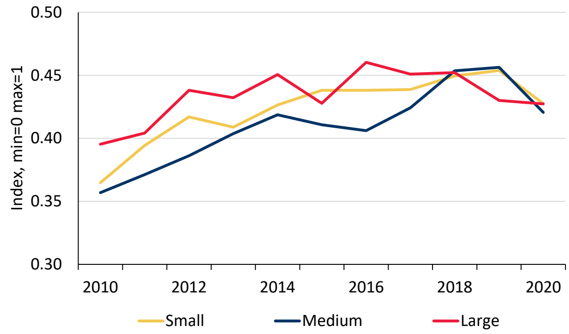

(A) 2010-2020

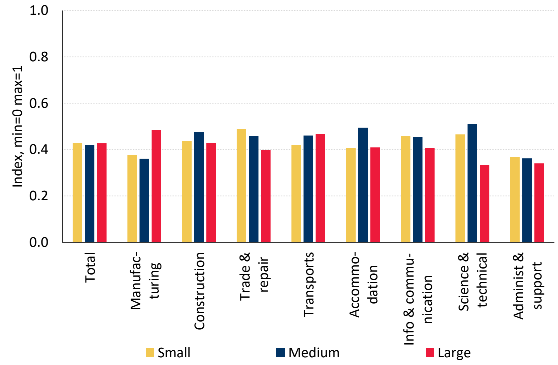

(B) Sectors, 2020

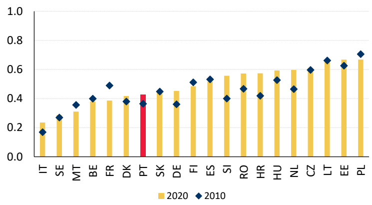

(C) Small firms

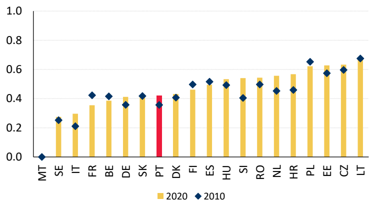

(D) Medium firms

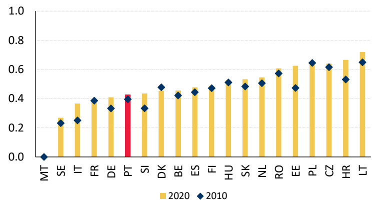

(E) Large firms

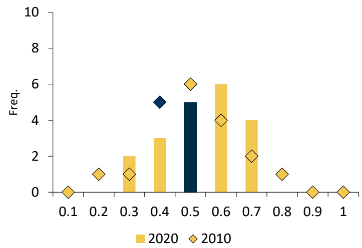

(F) Small firms

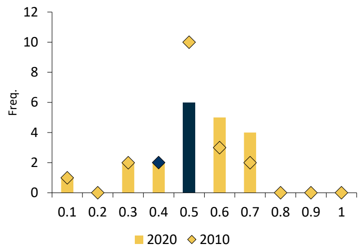

(G) Medium firms

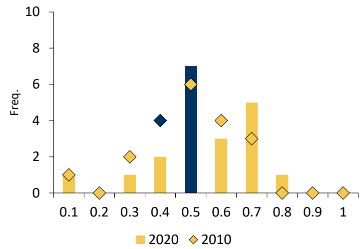

(H) Large firms

FIGURE 3: Equity ratio: Time, sector, rank across countries and distribution
Note: Blue bars in the histograms identify the bin where the median portuguese firm is.

benchmark the highest 75th percentile among countries, the distances between the median and this threshold are typically large due to the strong positive skewness of the distributions, a common feature in micro data variables. Panel B) of Figure 3 presents the indicator across sectors for the three size classes and the important result is the strong similarity among them. For all firm size classes and sectors, the interval of the indicator is broadly defined between 0.3 and 0.5.

As in the previous subsection, the second block of Figure 3 ranks countries according to the values of the indicator in 2020 for each firm size class. The median portuguese firm stands in the intermediate-low segment of the ranking in the three size classes. The corresponding histograms for 2020 presented in the last three panels of Figure 3 present a normal shaped distribution and median small portuguese firms stand near the mode, while medium and large firms are in their respective modal bins.

4.3. *Capital-labour ratio*

The comparatively low capital-labour ratios have been pointed out as a feature of portuguese firms. Panel A) of Figure 4 presents the path of the relative distance of this ratio in portuguese firms versus the highest 75th percentile of the distributions of EU firms. The striking result is the very low level of the our relative indicator, thus confirming the previous assessment. The path of the indicator for the period 2010-2020 does not show a relative progress in this situation, i.e., a reduction in the relative distance to the highest 75th percentile in the distributions of EU countries.

The analysis at the sector level in panel B) of Figure 4 shows that the relative high distance of the capital per worker in the median portuguese firm is pervasive across sectors, being especially high in sectors of Information & communication and Administrative & support services. There are also some differences in the indicator along firm-size classes in some sectors. The most salient one is the higher value for the large portuguese construction firms versus small and medium sized ones.

The ranking of the indicator for the three size classes along countries shows that the values are small in many of them, signaling again that the highest 75th percentile in the distributions of EU firms is quite distant from the median. The median portuguese firms are in an intermediate position in the rankings for the three firm-size classes. The three histograms in Figure 4 also show a positive skewness in the distribution of the indicator, portuguese small and medium firms are placed in the modal bin, while large ones are immediately below the mode.

4.4. *Labour productivity*

The (apparent) labour productivity is the first of two outcome variables we examine in this article. Panel A) of Figure 5 presents the path of the relative distance of labour productivity in the median portuguese firms in each size class relatively to the highest 75th percentile in the distribution of EU firms, in their respective size class, in the period 2010-2020. The figure shows a downward path for the indicator, i.e., an increasing distance to the best performers, for all size classes. This distance is larger for small firms, with the level of the indicator at 0.1 in 2020. This disappointing relative performance of labour productivity is compared along sectors in panel B) of Figure 5. The level of the indicator is low in all sectors, but particularly reduced in Trade & repair. In Construction the portuguese median large firm reaches a level of 0.2 in 2020, which is the maximum of the indicator across all sectors and size classes.

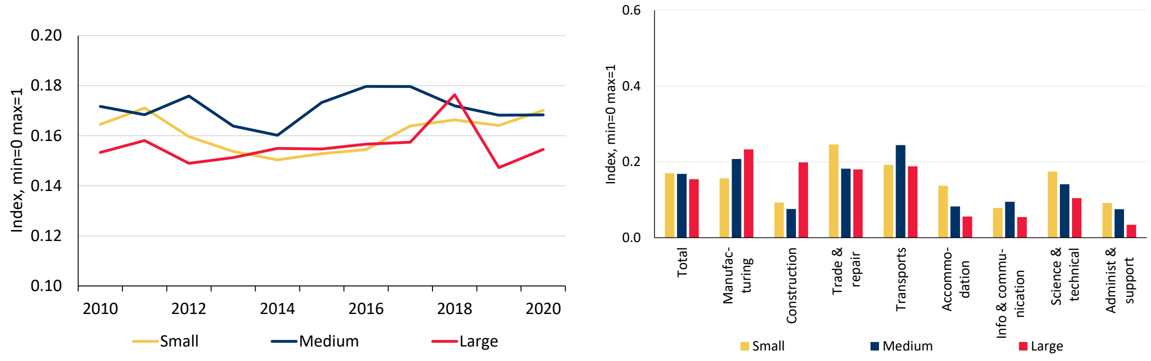

(A) 2010-2020 (B) Sectors, 2020

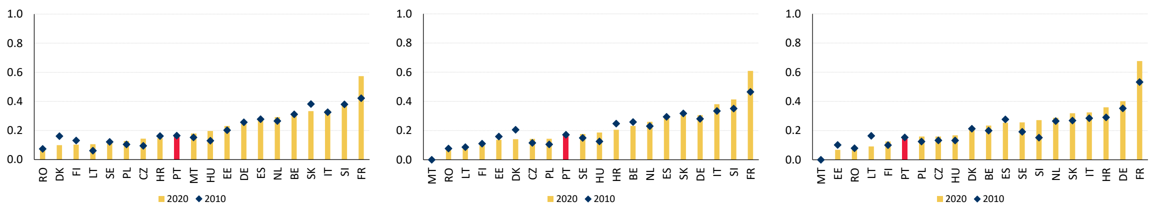

(C) Small firms (D) Medium firms (E) Large firms

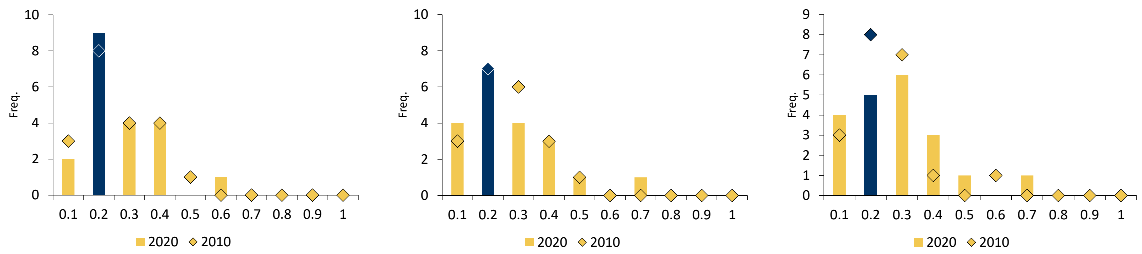

(F) Small firms (G) Medium firms (H) Large firms

FIGURE 4: Capital-labour ratio: Time, sector, rank across countries and distribution
Note: Blue bars in the histograms identify the bin where the median portuguese firm is.

The three panels in the second block of Figure 5 rank the situation of firms in the EU countries considered in 2020 and show that, although there are large distances relatively to the best performing firms, there are important differences across countries. Sweden, Belgium and France stand out in terms of lower distance of the median firms to the highest 75th percentile in the EU countries' distributions, while Portugal is in the lower part of the ranking, especially in the case of small firms. When these results are transposed to the position of portuguese median firms in the histograms of Figure 5, the position of portuguese firms in the group of countries with lower labour productivity, i.e, those in the lower bins of the histogram.

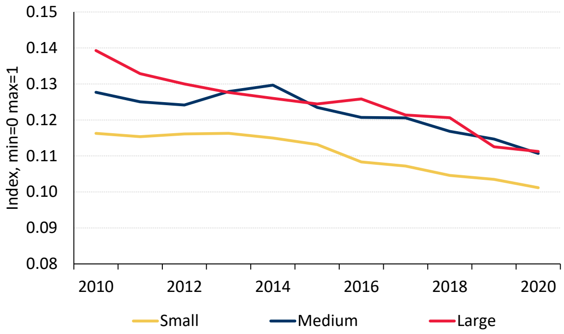

(A) 2010-2020

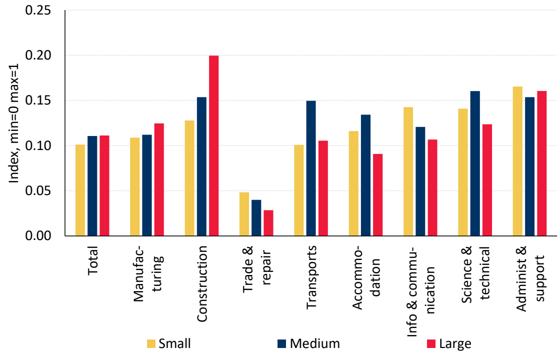

(B) Sectors, 2020

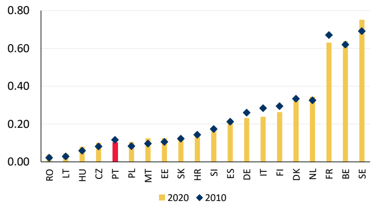

(C) Small firms

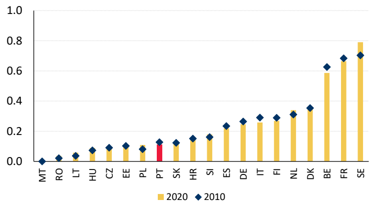

(D) Medium firms

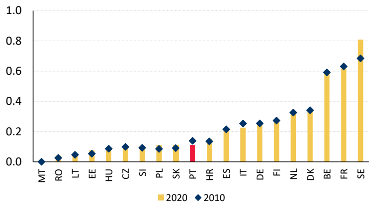

(E) Large firms

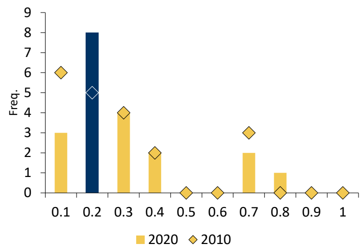

(F) Small firms

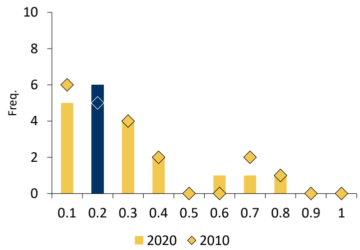

(G) Medium firms

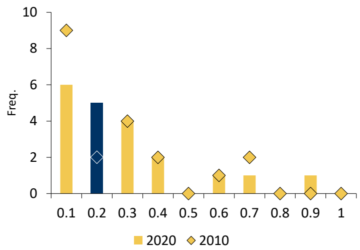

(H) Large firms

FIGURE 5: Labour productivity: Time, sector, rank across countries and distribution
Note: Blue bars in the histograms identify the bin where the median portuguese firm is.

4.5. *Return on assets*

The ROA is an indicator that tends to fluctuate along the economic cycle and sectoral shocks. This volatility is visible in panel A) of Figure 6. Nevertheless, beyond the yearly fluctuations, two facts emerge. The relative profitability of portuguese firms is low relatively to the firms in the highest 75th percentile of the EU distributions, with the largest firms performing better than the medium size firms and the latter performing relatively better than small ones.

In addition, Panel B) of Figure 6 presents the levels of the indicator in 2011 and 2020 for the sectors of activity in Portugal. The large firms in the sector of Information & communication and the small firms in the sector of Accommodation are those with a relatively smaller distance to the highest 75th percentile in their respective categories.

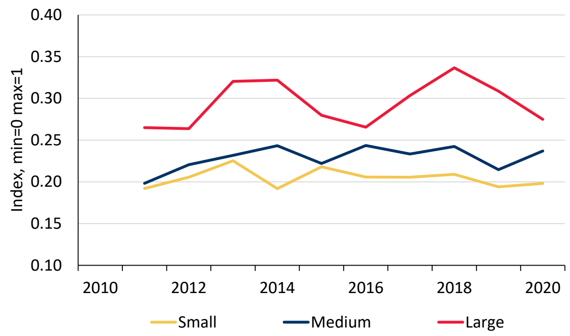

(A) 2010-2020

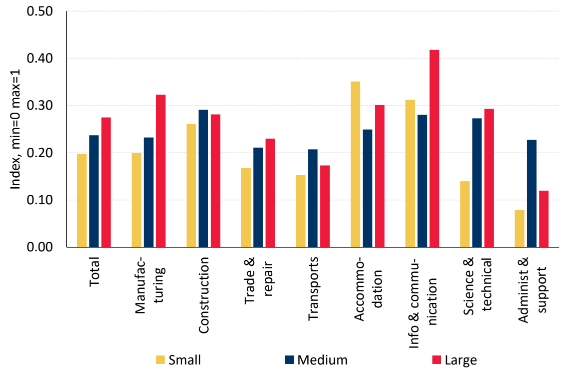

(B) Sectors, 2020

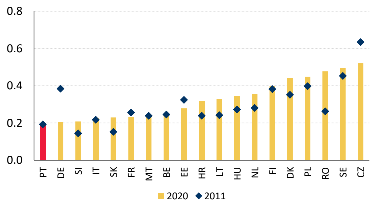

(C) Small firms

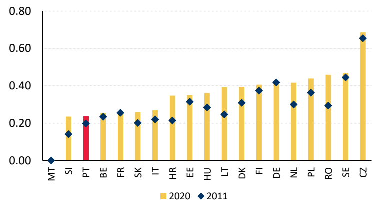

(D) Medium firms

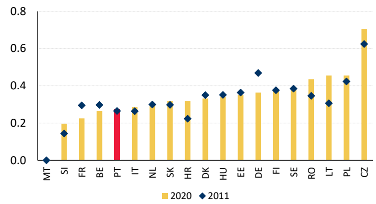

(E) Large firms

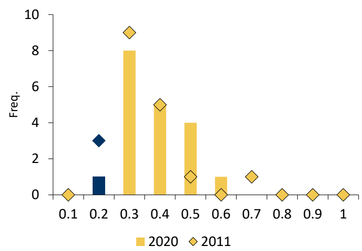

(F) Small firms

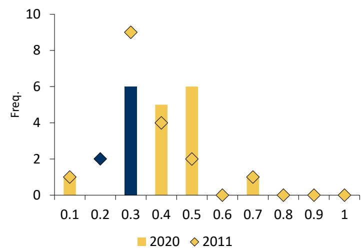

(G) Medium firms

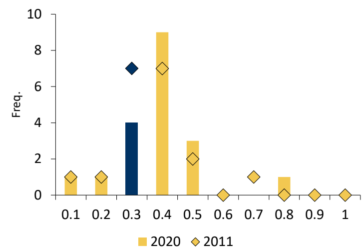

(H) Large firms

FIGURE 6: Return on assets: Time, sector, rank across countries and distribution
Note: Blue bars in the histograms identify the bin of the median portuguese firm.

The three panels of Figure 6 rank the 20 EU countries in our sample in 2020 for each firm size class. The sequence described in the previous paragraph repeats itself in terms of ranking. The position of median portuguese firms in the ranking is progressively worse for large, medium and small firms. Finally, when this information is presented in terms of the position in the histograms of the distribution of the indicator, the unfavorable position of portuguese small firms in terms of ROA is again visible.

5. Correlation and comparison with an aggregate data approach

In this section we correlate input and outcome variables. The objective is to assess if expected basic relationships between variables hold in the data and in our relative distance measure. In particular, we expect to identify a relationship between wages, capital per worker and labour productivity. In addition, we select the latter variable and compare the results of our micro-aggregated approach with a similar calculation based on aggregate official national accounts. The comparison of both layers of analysis is also a way to confirm results.

5.1. Correlation between input and outcome variables

In order to assess the firm-distribution data and the results presented in the previous sections we proceed with two exercises. Firstly, we plot scatter graphs between capital-labour ratios versus labour productivity, and capital-labour ratios versus average wages. Secondly, we run robust regressions to assess the correlation between our three input variables and each of the two outcome variables separately.

The six panels of Figure 7 present the scatter plots between capital-labour ratios versus labour productivity, and capital-labour ratios versus average wages in each country considered, separately for the three firm size classes. The dot for each country corresponds to the 2010-2020 average of the the indicator in each sector aggregated with its respective weight on GVA. As expected, the basic relationship between the two pairs of variables is consistently positive for each firm size class. The position of portuguese firms (dot in red) stands below the linear regression line, meaning that the relative labour productivity and the average wages are lower than those of countries with similar values for the capital-labour ratio. This adds to the narrative that the portuguese relative capital-labour ratio is low, as discussed in Subsection 4.3.

Moving beyond the scatter plots involves running a regression that controls for time, country, sector and firm-size specific effects. Table 5.1 presents the coefficients of two regressions that include the three input variables as regressors and either labour productivity or ROA as independent variables, also considering year fixed effects and combined sector-country-size fixed effects. Needless to say that these regressions simply provide a more solid statistical correlation than scatter graphs, and do not represent any causal relationship. In column (1), the coefficients for average wage and capital labour ratio are positive and statistically significant. This means that the proximity to the capital-labour ratio to the best performances is positively and significantly associated with the proximity to the labour productivity frontier, as expected in any production function. In addition, proximity to the average wage frontier is also found to be positively and significantly correlated with proximity to the labour productivity frontier. As for coefficients on the ROA equation, in column (2), the equity ratio is positive and statistically different from zero, while the coefficient for the capital labour ratio is negative.

5.2. Aggregated versus micro-aggregated labour productivity

In this section we provide an additional robustness analysis by comparing the labour productivity indicator that emerges from our micro-aggregated approach with an indicator computed in the same way, but using the aggregate labour productivity data published in the official national accounts. The latter data source is broader because it considers a larger number of sectors, including those associated with the activity

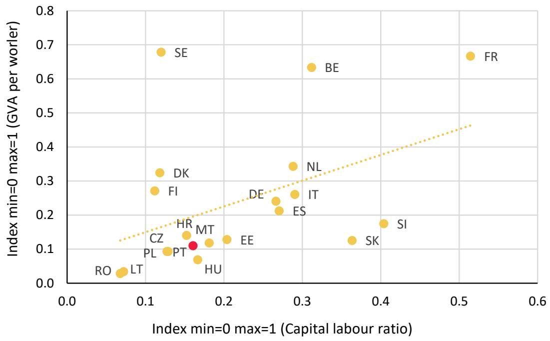

(A) Small firms

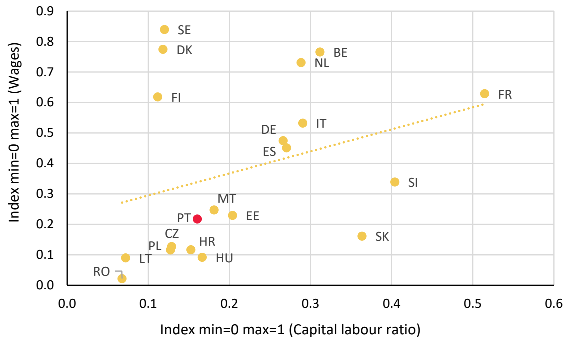

(B) Small firms

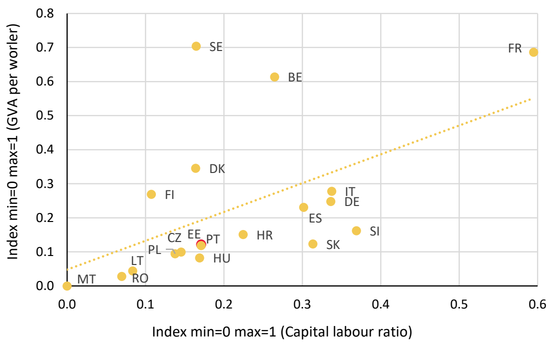

(C) Medium firms

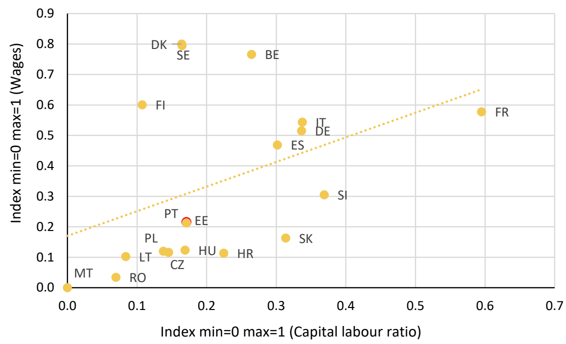

(D) Medium firms

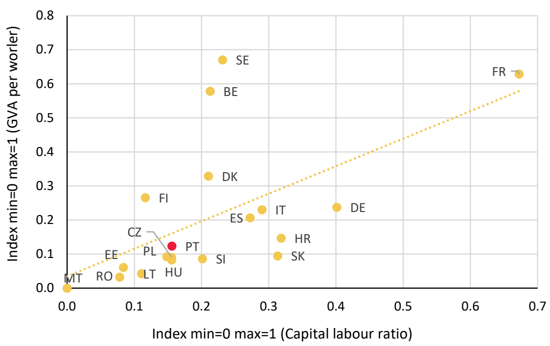

(E) Large firms

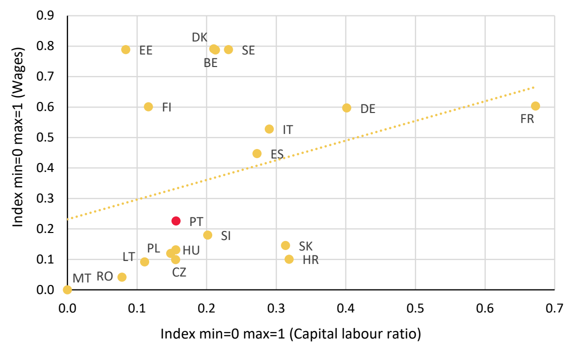

(F) Large firms

FIGURE 7: Labour productivity and average wages vs capital-labour ratio

Note: The indicator in the axes corresponds to Eq. 1. The line in each panel stands for the simple linear regression.

	(1) Labour prod.	(2) ROA
Average wage	0.277***	-0.0402
	(0.0284)	(0.0765)
Equity ratio	0.0131	0.112***
	(0.0150)	(0.0339)
Capital-labour ratio	0.0731***	-0.124***
	(0.0255)	(0.0393)
Constant	0.127***	0.340***
	(0.0125)	(0.0376)
Observations	5111	4774
R-squared	0.0924	0.0316
country x sector x size FE	✓	✓
year FE	✓	✓
Robust SE	✓	✓

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 2. Correlation between input and outcome variables

of general government, imputed rents, and other adjustments such as the informal economy. Nevertheless, the comparison is informative, mostly in what concerns the time path of the indicators.

Figure 8 presents the path of the relative labour productivity indicator, as defined in equation 1 for the period 2010-2020 with the one obtained with official national accounts, as presented in Amador *et al.* (2022). Unsurprisingly, the level of the indicator based on aggregate national accounts data is higher than the micro-aggregated one. In fact, the aggregate indicator includes all sectors of activity and setting the threshold for best performance at the country level is less demanding than setting it as the firm in the highest 75th percentile of the EU distributions, which are highly positively skewed.[1] Nevertheless, the path of the two indicators in the period 2010-2020 is downward. This conveys a robust message on the somewhat disappointing path of portuguese labour productivity, when measured as the distance to the best performances in the EU.

The three panels of Figure 9 compare the path of the micro-aggregated labour productivity indicator in each size class with the indicator based on aggregate national accounts data, using an index based in 2010. Although the share of each firm-size class is not the same in all sectors of the economy, the comparison is useful. The panels show that the aggregate indicator has decreased relatively more than those of the three types of firms, even if in case of large firms this decrease is more muted. The strong decrease in the aggregate indicator in 2020 should not be overstated because it corresponds to the first year of the COVID-19 pandemics, which had asymmetric and not synchronized effects across the EU countries.

1. In Amador *et al.* (2022) the best performer in each indicator (including labour productivity) is the EU country in the 90th percentile of the distribution of the respective variable.

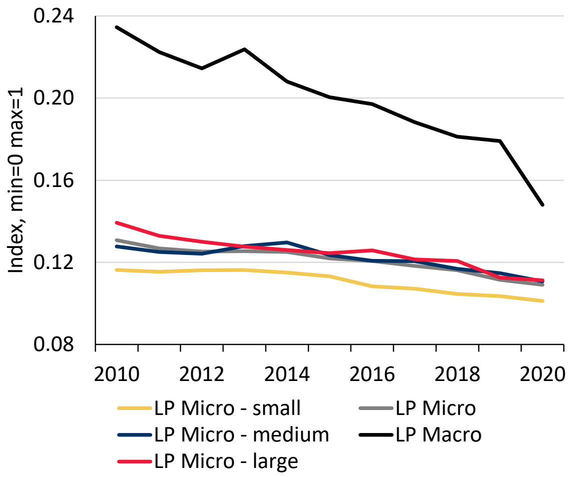

FIGURE 8: Micro-aggregated versus aggregate labour productivity indicators
Note: The macro indicator corresponds to that of Amador *et al.* (2022).

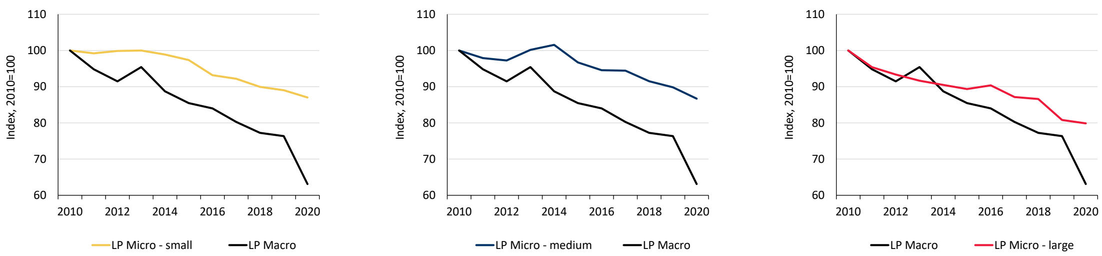

(A) Small firms (B) Medium firms (C) Large firms

FIGURE 9: Micro-aggregated versus aggregate labour productivity indicators (2010=100)
Note: The macro indicator corresponds to that of Amador *et al.* (2022).

6. Final remarks

This article takes a micro-aggregated approach to describe the relative position of portuguese firms versus the best performing ones in the EU, defined as those in the highest 75th percentile of the distributions of the relevant indicator in the EU countries. A set of three input variables (average wages, equity ratio and capital-labour ratio) and two outcome variables (labour productivity and ROA) are studied, along time, sector and firm size class. This micro-based and cross-country comparative approach significantly improves upon single country firm-level studies, and cross country aggregate studies. In this type of approach it is possible to unveil differences among firms of similar size in the same sectors across different EU countries.

Overall, the median portuguese firm in each size class and sector exhibits a significant performance gap relative to the best performing firm in the EU in the corresponding category. This distance is acute in the capital-labour ratio, average wages and labour productivity, which are three intrinsically related variables. The positive correlation between these variables clearly emerges in our cross-country firm-level data. The path of the relative indicators in the period 2010-2020 shows some convergence toward the best performing EU firms in the equity ratio and widening gaps in terms of average wages and labour productivity. The somewhat disappointing path of labour productivity is consistent with results obtained using a similar methodology but based on national accounts aggregate data.

The agenda for future research on productivity in Portugal and the EU is too long to be listed here and it is dependent on the availability of granular comparable data. There is an increasing need to establish comparisons across firms in different countries in order to assess the impact of firms' characteristics on productivity over and above the national specificities. The set of relevant firms' characteristics is large, but the dimensions related to skills, R&D and adoption of new digital technologies are likely to be key. Only with a deeper knowledge of firm-level productivity drivers and related market failures can public policies and investment packages be efficiently allocated. The ongoing debate about industrial policies and the subsequent design of such interventions should take these concerns aboard, instead of going towards a blanket support for specific sectors or types of firms.

References

Adilbish, Oyun, Diego Cerdeiro, Romain Duval, Gee Hong, Luca Mazzone, Lorenzo Rotunno, Hasan Toprak, and Maryam Vaziri (2025). "Europe's Productivity Weakness: Firm-Level Roots and Remedies." *IMF Working Papers*, 2025(040).

Amador, João, Ana Fernandes, and Guida Nogueira (2022). "The competitiveness of the Portuguese economy: A view from a composite indicator." *Economic Bulletin and Financial Stability Report Articles and Banco de Portugal Economic Studies*.

Arnold, Jens Matthias and Natália Barbosa (2015). "Structural policies and productivity: Evidence from Portuguese firms." OECD Economics Department Working Papers 1259, OECD Publishing.

Banco de Portugal (2019). "Labour productivity in Portugal over the past decade: A firm-level approach." Economics bulletin, special issue, Banco de Portugal.

Bartelsman, Eric J, Filippo Di Mauro, Sergio Inferrera, Marco Matani, Ugo Panizza, and Michael Polder (2023). "Microdata for Economic Research in Europe: Challenges and Proposals." CEPR Discussion Papers 18640, C.E.P.R. Discussion Papers.

Bergeaud, A. (2024). "The past, present and future of European productivity." Report, ECB Forum on Central Banking.

CompNet (2014). "Micro-based evidence of EU competitiveness: The CompNet database." Working Paper Research 253, National Bank of Belgium.

Connor, Tim (2003). "Managing for competitiveness: A proposed model for managerial focus." *Strategic Change*, 12, 195–207.

CPP (2019). "A Produtividade da Economia Portuguesa." Policy research working paper series, Conselho para a Produtividade.

Draghi, Mario (2024). "The future of European competitiveness: A competitiveness strategy for Europe." Report, European Commission.

Dunn, Malcolm H. (1994). “Do Nations Compete Economically? A Critical Comment on Prof. Krugmanâs Essay âCompetitiveness: A Dangerous Obsessionâ.” *Intereconomics*, 29(6), 303–308.

European Commission (2025). “An EU Compass to regain competitiveness and secure sustainable prosperity.”

Gomis, Roger M. and Sameer Khatiwada (2017). “Debt and productivity: Evidence from rm-level data.” IHEID Working Papers 04-2017, Economics Section, The Graduate Institute of International Studies.

Krugman, Paul (1994). “Competitiveness: A Dangerous Obsession.” *Foreign Affairs*, 73(2), 28–44.

Krugman, Paul (1996). “Making Sense of the Competitiveness Debate.” *Oxford Review of Economic Policy*, 12(3), 17–25.

Letta, Enrico (2024). “Much More Than a Market.” Report, European Commission.

Martins, Fernando, Hugo Reis, Manuel Coutinho Pereira, Mário Lourenço, and Cloé Magalhães (2022). “An assessment of companiesâ competitiveness in Portugal and in some European countries.” *Economic Bulletin and Financial Stability Report Articles and Banco de Portugal Economic Studies*.

Peneder, Michael (2017). “Competitiveness and industrial policy: from rationalities of failure towards the ability to evolve.” *Cambridge Journal of Economics*, 41, 829–858.

Solow, Robert M. (1957). “Technical Change and the Aggregate Production Function.” *The Review of Economics and Statistics*, 39(3), 312–320.

Non-technical summary

October 2025

The vulnerability channel: assessing the impact of financial conditions on the output gap

Ivan De Lorenzo Buratta, Diana Lima

The article explores how financial vulnerabilities affect macroeconomic outcomes, particularly focusing on the asymmetry in the distribution of output gap changes in Portugal. It builds on the "Growth-at-Risk" (GaR) literature, which highlights that financial conditions disproportionately affect downside risks (tail risks) in economic growth.

Traditional New Keynesian (NK) models assume linearity and constant volatility, failing to capture real-world asymmetries. Using quantile regressions, we show that in Portugal, as in the U.S. (Adrian et al., 2019 and Adrian et al., 2020), tight financial conditions primarily worsen the lower tail of the output gap change distribution. This means bad economic outcomes become more likely under financial stress. Inflation and past output gap changes have minimal impact, reinforcing that financial stress is the main driver of asymmetry.

The New-Keynesian Vulnerability (NKV) model extends the standard structural macroeconomic framework, known as New Keynesian model, by incorporating a financial vulnerability channel, where volatility depends on past financial conditions and output gaps. This structure captures key features of the Portuguese economy, such as non-constant volatility and skewness in output gap changes. The model, estimated with Portuguese data (1994–2019), successfully reproduces empirical findings (Figure 1). First, volatility in the output gap change correlates negatively with its conditional median. Second, the left-hand tail (corresponding to severe downturns) responds strongly to worsening financial conditions. Third, simulations show that intertemporal risk-taking increases during periods of low volatility, leading to future instability.

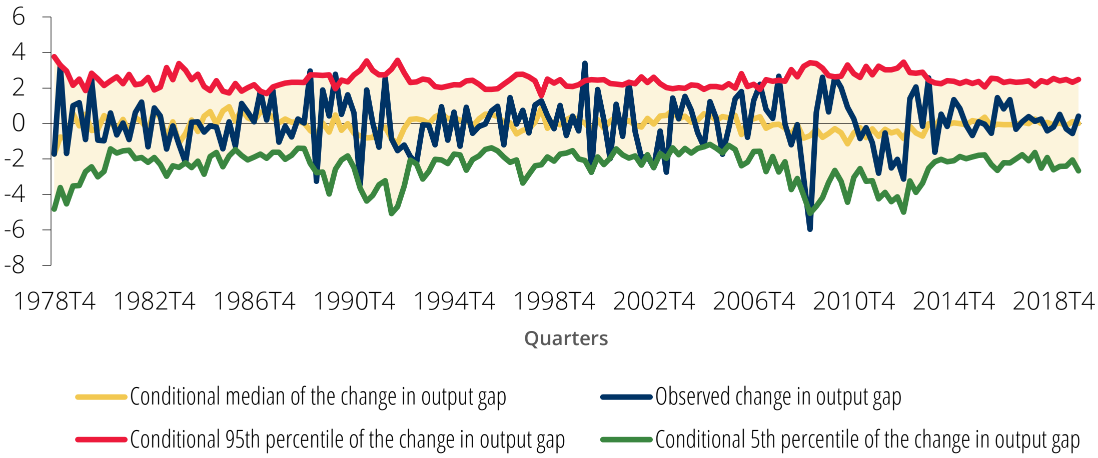

FIGURE 1: Forecast of the changes in the output gap | percentage

Note: results from one-quarter-ahead quantile regressions that include the current output gap change, inflation and the Country-Level Index of Financial Stress (CLIFS) as regressors.
Source: Banco de Portugal, authors' calculations.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. XI, No. 4

The vulnerability channel: assessing the impact of financial conditions on the output gap

Ivan De Lorenzo Buratta
Prometeia

Diana Lima
Banco de Portugal

October 2025

Abstract
This study estimates a New Keynesian model with a financial vulnerability channel (NKV model) for the Portuguese economy, building on the framework proposed by Adrian *et al.* (2020). Motivated by the "Growth-at-Risk" (GaR) literature, which emphasizes the importance of non-linearities in macro-financial dynamics, the model introduces state-dependent volatility to capture asymmetries in the distribution of output gap realizations. Empirical evidence shows that financial conditions significantly influence the lower tail of GDP growth, with deteriorating conditions increasing downside risks. Our results indicate that the NKV model successfully replicates these non-linear features for Portugal: financial stress amplifies the probability of extreme negative outcomes, while low current volatility can trigger future risk-taking and greater macroeconomic instability. The model highlights a key intertemporal risk trade-off for policymakers and supports the case for taking into account the interactions between macroprudential and monetary policy in the stabilization of financial conditions. (JEL: E12, E23, E37)

1. Introduction

The recent boom of the "Growth-at-risk" literature stems from the need to integrate non-linearities in macroeconomic and financial dynamics, i.e., explaining what drives the asymmetry in the distribution of some macroeconomic variables' realizations (e.g., GDP growth, house prices or credit growth). The empirical literature (among others, Aikman *et al.* (2019), Adrian *et al.* (2019)) agrees on defining financial vulnerability and aggregate risk as key drivers of "tail risks" of GDP growth and other macroeconomic variables – that is, how severe a downturn of these variables could become in the case of extreme events related to financial conditions and aggregate drivers. Adrian *et al.* (2019) show that financial conditions play a critical role for forecasting the lower tail of US GDP growth. Deteriorating financial conditions are

Acknowledgements: The authors thank the participants of the European Central Bank WGEM-WGF Expert Group on Macro-at-Risk, as well as the editor, an anonymous referee, Ana Cristina Leal, Ana Lacerda, Duarte Maia and Fátima Silva for their comments. Special thanks to Leonardo Urrutia, for sharing the codes used in the estimation of the model. The analyses, opinions, and conclusions expressed in this study are those of the authors and do not necessarily reflect those of the Banco de Portugal or the Eurosystem.

E-mail: ivan.delorenzoburatta@prometeia.com; dlima@bportugal.pt

associated with an increase in the conditional volatility and a decrease in the conditional mean of GDP growth. Therefore, financial conditions have significant predictive power for future GDP vulnerability, whereas economic conditions are not as informative for predicting tail outcomes.

This argument also applies to Portugal, since, as showed in De Lorenzo Buratta *et al.* (2022a), De Lorenzo Buratta *et al.* (2022b), and Passinhas and Pereira (2023), financial stress and cyclical systemic risk measures have heterogeneous effects – across time and percentiles – on Portuguese GDP growth, credit growth, house price growth, and banks' profitability.

Structural macroeconomic models such as New Keynesian (NK) models – macroeconomic models that assume rational expectations of agents, imperfect competition, and price "stickiness" – have great potential for policy analysis and are frequently used by central banks to assess the impacts of monetary and macroprudential policies. Such models are also useful for assessing financial stability implications of macroprudential policies and its interplay with monetary policy. However, despite the empirical evidence, most macroeconomic models of this sort still rely on linearization approaches and therefore cannot generate asymmetries as documented in the empirical growth-at-risk (GaR) literature. Even for more versatile models such as linear ones with occasionally binding constraints, it is challenging to generate asymmetries such as the ones observed in the data. Some recent papers started exploring these empirical non-linearities in standard NK models, trying to enrich the framework while guaranteeing tractability. These papers build on earlier work that investigates the role of time-varying uncertainty in Dynamic Stochastic General Equilibrium (DSGE) models (Christiano *et al.* (2014); Cesa-Bianchi and Fernandez-Corugedo (2018)).Aikman *et al.* (2023), for example, replicate the non-linear features of GDP growth by defining a lower bound on interest rates, a credit crunch that triggers when bank capital depletes, and a deleveraging of borrowers that activates when their debt service burden is excessive.

In addition, any standard NK model also postulates the constancy of macroeconomic variables' volatility, although recent studies have suggested that this might be an assumption that is not underpinned by data.

In this study, we follow the work of Adrian *et al.* (2020), which suggests a simple model to understand the mechanisms of financial vulnerabilities and their impact on the output gap dynamics. Growth-at-risk dynamics are included in an otherwise standard NK model by introducing a financial vulnerability channel, embedding nonlinear volatility that leads to state-dependent second and higher moments. This "New Keynesian Model with Vulnerability" – henceforth the NKV model – captures a rich risk structure, allowing to show that the one-quarter-ahead estimated distribution of the output gap change is highly asymmetric for the United States. They argue that this result is driven by financial conditions, impacting the left-hand tail of the output gap change distribution – corresponding to the lowest observations – while leaving the right-hand tail – corresponding to the highest observations – mostly unaffected. The NKV model estimated for the US economy reproduces key GaR features that are absent from fully linear structural models, such as the asymmetry in conditional output gap percentiles or the skewness and kurtosis in its distribution. This type of kurtosis-related nonlinearity is relevant for the study of the link of tail risks to output stemming from poor financial conditions. Nevertheless, Carriero *et al.* (2024) find some evidence that tail risks can be captured not only with models that allow asymmetries in conditional distributions, but also with a Bayesian Vector Autoregressions approaches with conventional stochastic volatility that yields symmetric conditional distributions.

The goal of this study is to estimate the NKV model with data for Portugal and evaluate the capacity of the model to reproduce the empirical features of the Portuguese output gap distribution. The inclusion of a financial vulnerability channel in a standard NK model illustrates the importance of current and expected financial conditions for the future realizations of the output gap in Portugal, and possibly shed light on the intertemporal costs and benefits for macroprudential policy action in stabilising financial conditions.

The next sections are structured as follows. Section 2 estimates the empirical features of the output gap change distribution for Portugal. Section 3 presents the NKV model and the estimation results. Section 4 describes a sensitivity analysis exercise, which explores the impact on the output gap of more volatile financial conditions. Section 5 concludes.

2. Output gap change: empirical features

Any standard NK model assumes the constancy of conditional second moments. As showed in Adrian *et al.* (2020), this feature is rejected by data, as the conditional median and volatility of the changes in the output gap are negatively correlated. We start by showing in this section that the empirical results in Adrian *et al.* (2020) for the United States also hold for Portugal. This is an important step to motivate the estimation of the NKV model with Portuguese data.

For this purpose, we estimate the one-quarter-ahead output gap change distribution using quantile regressions,[1] where the current output gap change, inflation, and the Country-Level Index of Financial Stress (CLIFS) are used as independent variables. We estimate the marginal effect of each explanatory variable on the evolution of the 1-quarter-ahead dependent variable, for each percentile q between 1% and 99% in steps of 1 percentage point, as follows:

$$\hat{Q}_{\triangle y_t^{gap}}(q|\triangle y_{t-1}^{gap}, infl_{t-1}, CLIFS_{t-1}) = \hat{\alpha}^q + \hat{\beta}^q_{\triangle y^{gap}} \triangle y_{t-1}^{gap} + \hat{\beta}^q_{infl} infl_{t-1} + \hat{\beta}^q_{CLIFS} CLIFS_{t-1}, \quad (1)$$

where $\triangle y_t^{gap}$ is the dependent variable, $\triangle y_{(t-1)}^{gap}, infl_{(t-1)}, CLIFS_{(t-1)}$ are the explanatory variables, $\hat{\alpha}^q$ represents the model estimated constant and $\hat{\beta}^q_{\triangle y^{gap}}$, $\hat{\beta}^q_{infl}$, $\hat{\beta}^q_{CLIFS}$ are parameters that represent the estimated marginal effect of each explanatory variables on the dependent variable.

1. Brownlees and Souza (2021) conduct an out-of-sample empirical evaluation of Growth-at-Risk (GaR) predictions for 24 OECD countries, comparing forecasts from quantile regression models and GARCH models and find that GARCH models are generally more accurate than quantile regression models.

We assume the estimated 5^{th} and 95^{th} percentiles to be the thresholds representing the lowest and highest realizations of the one-quarter-ahead output gap change, respectively.

Figure 1 shows the estimated one-quarter-ahead changes in the output gap. According to these results, the lowest realizations of the one-quarter-ahead output gap change distribution exhibits 2.17 times the standard deviation of the highest realizations. This difference in volatility causes the asymmetry of the estimated distribution of the output gap change.

Figure 2 indicates that the difference in volatility between the 5^{th} and 95^{th} percentiles of the estimated distribution of the output gap change is fully explained by the CLIFS, which is a measure of financial risk materialisation and a proxy of financial conditions dynamics in the regression. When the variable CLIFS increases by one-standard deviation, indicating a tightening of financial conditions, the estimated marginal effect on the lowest values of the output gap change decrease by 15 p.p. while it is not statistically significant in higher values. This result implies that high risk materialisation today, corresponding to tight financial conditions, has a more negative effect on the lowest values of the output gap change in one quarter compared to the rest of its values.

The other variables in the regression are not significant for any percentile. In other words, the right-hand tail of the distribution of the output gap stays roughly constant, while the left-hand tail varies with financial conditions. We can thus conclude that financial conditions, by driving the difference in volatility between the tails of the distribution of the output gap change, are responsible for the non-linear developments of the output gap change. An increasing volatility of the output gap change corresponds to higher probability of observing extreme negative output gap change realizations.

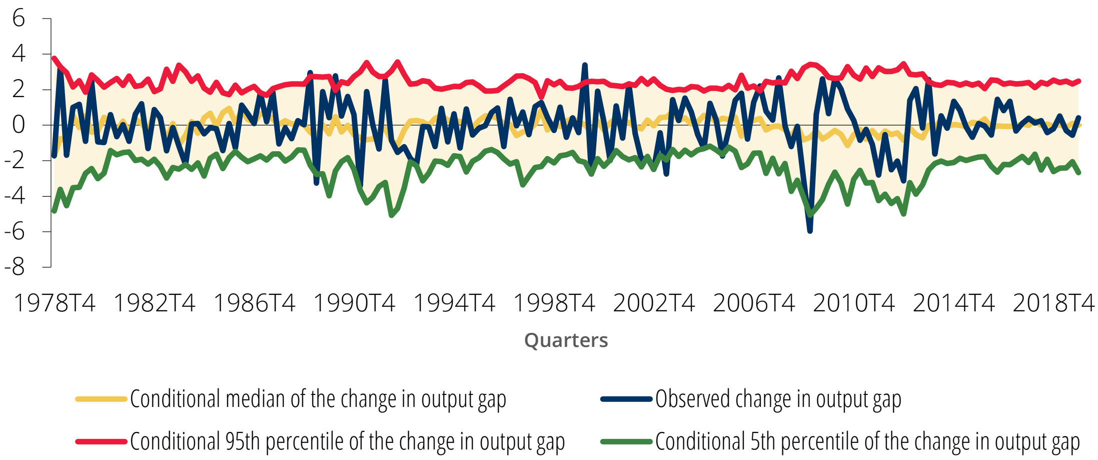

FIGURE 1: Forecast of the changes in the output gap | percentage

Note: results from one-quarter-ahead quantile regressions that include the current output gap change, inflation and the Country-Level Index of Financial Stress (CLIFS) as regressors.
Source: Banco de Portugal, authors' calculations.

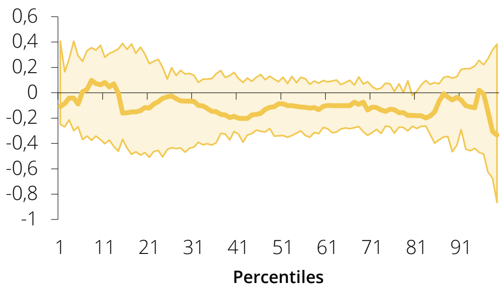

(A) Output gap change (lagged)

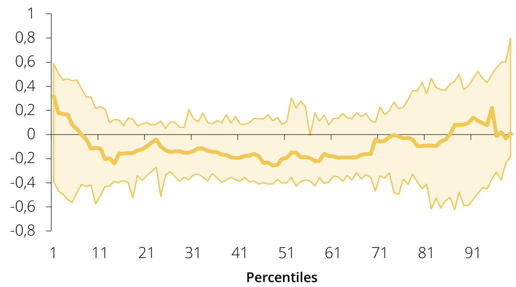

(B) Inflation

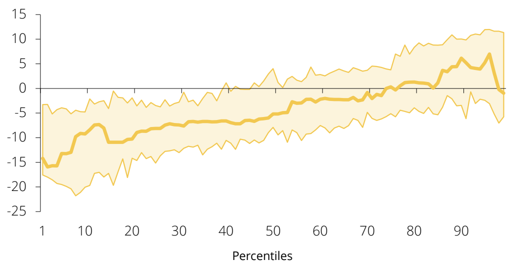

(C) Country-Level Index of Financial Stress

FIGURE 2: Marginal contributions to the output gap change | percentage points

Note: results from one-quarter-ahead quantile regressions that include the current output gap change, inflation and the Country-Level Index of Financial Stress (CLIFS) as regressors. The shaded areas stand for 95% confidence intervals obtained using bootstrapping (xy-pair method) according to Davino et al. (2013). The estimated marginal effects are conditional on a one standard deviation increase holding constant all other regressors.
Source: Banco de Portugal, authors' calculations.

Figure 3 shows that the volatility of the estimated distribution of the output gap change – measured as the difference between the estimated 95th and 5th percentiles – is negatively correlated with its conditional median. In other words, low realizations of the output gap change are associated with periods in which its extreme negative realizations are more likely to occur. The empirical results suggest overall that the output gap change volatility is not constant also for Portugal and that, similarly to the results from Adrian *et al.* (2020), this is due to financial conditions.

3. The model, data, and estimation

The core novelty of the Adrian *et al.* (2020)'s paper is the specification of the non-constant volatility of the output gap (i.e., the vulnerability channel), which is modeled to depend on past financial conditions and past values of the output gap. These modeling choices are justified by the stylized facts reported in Section 2, where we show that financial conditions today negatively affect the output gap change volatility in one quarter, and that there is a (negative) correlation between the output gap change volatility and its conditional median.

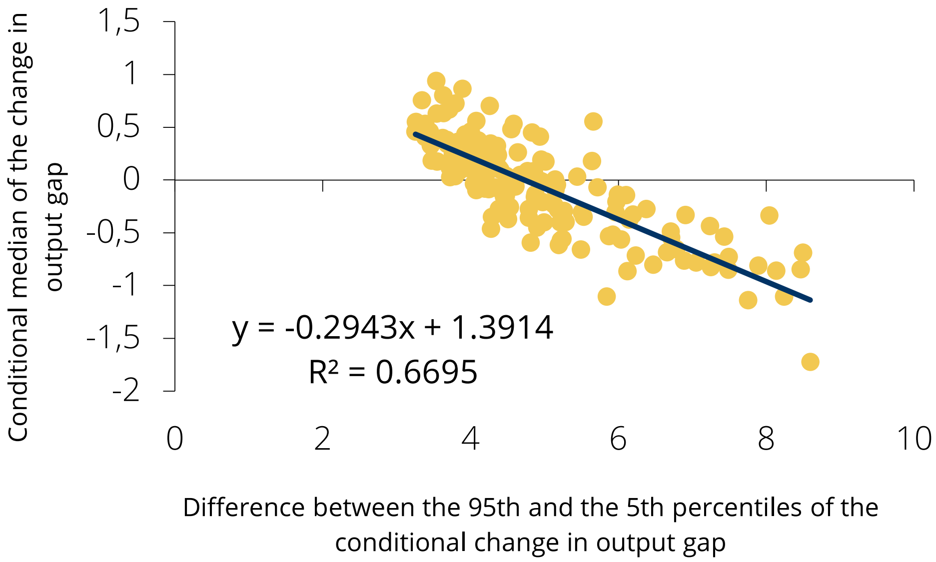

FIGURE 3: Correlation of the estimated one-quarter-ahead output gap change with its volatility | percentage

Note: results from one-quarter-ahead quantile regressions that include the current output gap change, inflation and the Country-Level Index of Financial Stress (CLIFS) as regressors.
Source: Banco de Portugal, authors' calculations.

The NKV model is an extension of a standard NK model that can be reduced to a three-equation system. The first equation is an IS curve, which models the combinations of the output gap and the interest rate that guarantee the general equilibrium in the real sector of the economy. The IS curve also includes a "financial accelerator" term, η_t (equation 2), to account for the well-known amplification effects of financial developments on economic realizations. Loose (tight) financial conditions have a positive (negative) impact on economic growth. In addition, the IS curve includes an extra wedge, $V(X_t)\varepsilon_t{}^{y_{gap}}$, to ensure that the variance of the shock to the equation is conditionally heteroskedasticity. While $\varepsilon_t^{y_{gap}}$ is i.i.d. $N(0, \sigma_y^2)$, $V(X_t)$ (in equation 3) represents vulnerability to exogenous shocks, which depends on past values of financial conditions and the output gap, varying with past state variables.

The fourth equation is a Phillips curve, which summarizes the positive relationship between inflation and the output gap. The last equation is a Taylor Rule, defining a standard monetary policy where the policy interest rate, i_t, depends on inflation and economic growth.[2] The authors also assume current financial conditions to be dependent on contemporaneous and expected output gap, with current and expected positive economic developments being associated with looser financial conditions today.

2. Papadamou *et al.* (2018) show that a single Taylor Rule does not fit all members of the euro area. We assume, however, that Portuguese economic activity responds to common euro area monetary policy shocks in a manner similar to the average member state. By contrast, macroprudential policy in Portugal is explicitly targeted at country-specific financial stability risks.

In the model, financial conditions eta_t depend indirectly on the interest rate i_t through the IS curve, and this feature allows the *risk-taking channel* of monetary policy. This vulnerability feature enables intertemporal risk-taking dynamics, with low current economic volatility triggering higher leveraging and causing higher expected future volatility in the output gap.

$$y_t^{gap} = E_t y_{t+1}^{gap} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1}\right) - \gamma_\eta \eta_t - V(X_t)\,\varepsilon_t^{y_{gap}} \tag{1}$$

$$\eta_t \equiv \lambda_\eta \eta_{t-1} + \lambda_{\eta\eta}\eta_{t-2} - \theta_y y_t^{gap} - \theta_\eta E_t y_{t+1}^{gap} \tag{2}$$

$$V(X_t) = v_c + \varrho_{\eta_1}\eta_{t-1} + \varrho_{\eta_2}\eta_{t-2} + \varrho_{gap} y_{t-1}^{gap} \tag{3}$$

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t^{gap} \tag{4}$$

$$i_t = \varphi^\pi \pi_t + \varphi^y y_t^{gap} + \nu_t \tag{5}$$

The model is fitted to Portuguese data (Figure 4) according to the existing literature on the estimation of DSGE models (Smets and Wouters (2003, 2007); An and Schorfheide (2007); Fernández-Villaverde and Rubio-Ramírez (2007).[3]

First, we estimate a subset of parameters for a linear version of the model, using pre-set standard values for the rest of them (chapter 3 of Galí (2008) textbook). Second, we estimate the non-linear parameters $(v_c, \varrho_{\eta_1}, \varrho_{\eta_2}, \varrho_{gap})$ matching some Portuguese empirical moments.[4] The following 1994-2019 time series are used as observables in the estimation (Figure 4): labour (hours), the Country-Level Index of Financial Stress (CLIFS), which is an indicator that comprises six, mainly market-based, financial stress measures that capture three financial market segments: equity markets, bond markets and foreign exchange markets (Duprey *et al.* (2015)), the euro area shadow short-term interest rate (Krippner, 2019), year-on-year core inflation, the output gap, the output gap change, the expected output gap change, and the estimated volatility of the output gap change.[5] The expected output gap change and the estimated volatility of the output gap change are obtained from the one-quarter-ahead quantile regression discussed in Section 2. Table 1 of the Appendix presents the estimated parameters.

---

3. This work was developed under the scope of the European Central Bank WGEM-WGF Expert Group on Macro-at-Risk, published as an ECB Occasional Paper "Using structural models to understand macroeconomic tail risks". Special thanks to Leonardo Urrutia from Leipzig University (Institute for Theoretical Economics - Macroeconomics) for providing the codes and help regarding the estimation. This section follows his methodological approach.

4. See Montes-Galdón *et al.* (2024) for details.

5. We experimented the use of alternative proxies for financial conditions, namely the Systemic Risk Index (Lang *et al.* (2019)) and the Composite Indicator of Systemic Stress (https://data.ecb.europa.eu/data/datasets/CISS/data-information). The Systemic Risk Index performs worse in our setting because it mainly affects Growth-at-Risk in the medium to long term for Portugal (see De Lorenzo Buratta *et al.* (2022a,b)), whereas our model is designed to capture one-quarter-ahead dynamics. The Composite Indicator of Systemic Stress is very similar to CLIFS, except in the most recent years, which we exclude to avoid contamination from COVID-19 quarters. Moreover, this series only begins in 1999, while CLIFS provides a longer history.

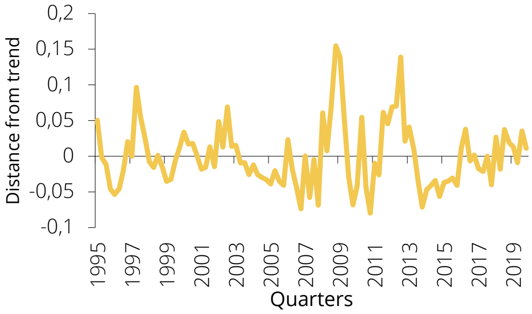

(A) Fin. Vulnerability (CLIFS)

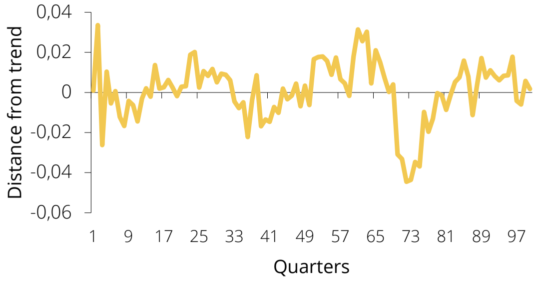

(B) Labor (Hours)

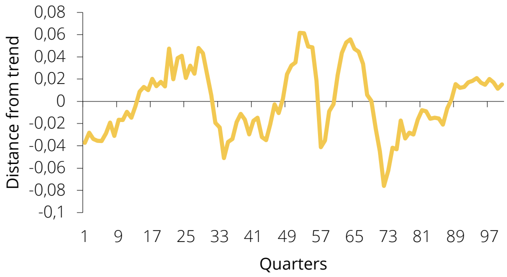

(C) Output gap

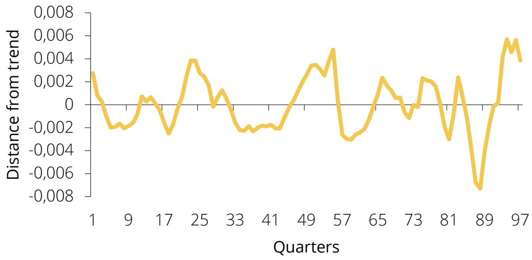

(D) SSR interest rate EU

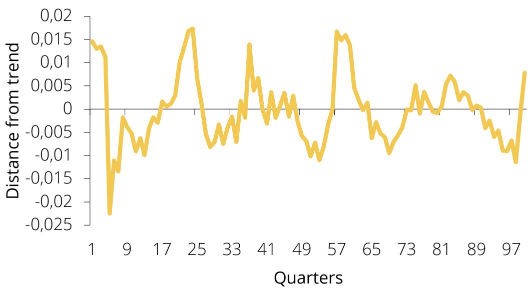

(E) Core inflation (OYA)

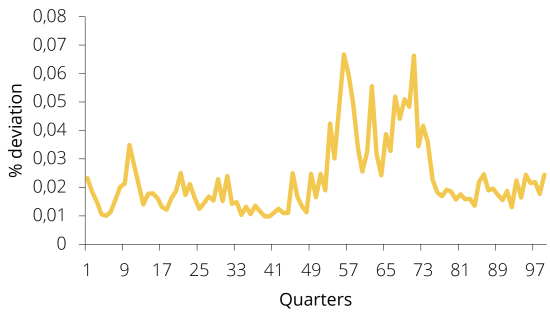

(F) *Conditional volatility of output gap change

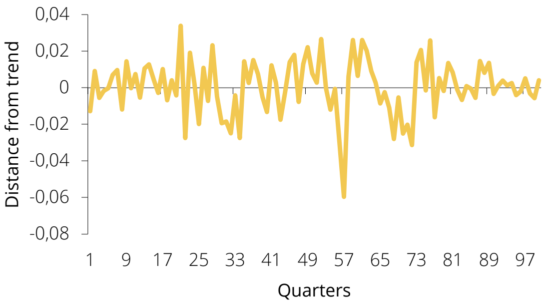

(G) Change in the output gap

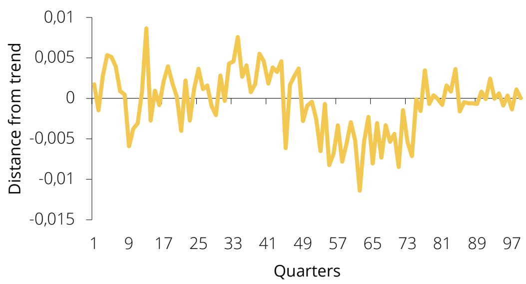

(H) *Conditional median of output gap change

FIGURE 4: Portuguese data series used as observables in the DSGE model estimation

Series are detrended/standardized. Variables with * are obtained from the quantile regression described in Section 2.

Source: Banco de Portugal, authors' calculations.

To compare the empirical results in Section 2 with those simulated by the model, we compute the median, the 5th and the 95th percentiles of the distribution of the output gap change, obtained by producing 200-quarter responses of the model to random supply-side, monetary-policy, and financial-conditions shocks.[6]

Figure 5 shows that the standard deviation of the 5th percentile of the output gap change is 2.12 times the standard deviation of the 95th percentile.

Figure 6 shows that the volatility of the estimated distribution of the output gap change is negatively correlated with its conditional median. The NKV model estimated for Portugal is thus able to replicate the stylized facts reported in Section 2. We conclude that simulation results obtained from the NKV model for Portugal match well the empirical evidence from the quantile regressions.

The non-linearity in the realizations of the output gap is driven by the time-varying nature of volatility in the model. With the purpose of illustrating the importance of current and expected financial conditions for the future realizations of the output gap, we consider our baseline model specification with non-constant volatility (i.e. with the vulnerability channel) and a specification where the volatility of the output gap is kept constant (without the vulnerability channel). We compute a 10000-period simulation of the two specifications, where state variables are randomly buffeted with supply-side, monetary-policy, and financial-conditions shocks.

Figure 7 indicates that the simulated distributions of the output gap and the output gap change obtained with the non-constant volatility specification are more asymmetric compared to the constant volatility specification, signalling that it is precisely the vulnerability channel of the model that is responsible for the non-linear realizations of the output gap. The more frequent positive values for the output gap change and the more frequent negative values for the output gap in the non-constant specification indicate that the economy changes more rapidly and is more likely to experience recessions when the financial conditions drive the output gap volatility.

4. Sensitivity analysis: the impact on the output gap of more volatile financial conditions

In this study, we are also interested in observing the consequences of having more/less volatile financial conditions in the model, to get some insight into the impact of policies that may reduce their volatility. Given the introduction of the vulnerability channel in a standard NK model, we can model a scenario with more volatile financial conditions. We compute impulse response functions – with random shocks perturbating the model over a 10000-period simulation as in the previous exercise – for the variables of the baseline specification with non-constant volatility and for an alternative specification

6. In the simulations, financial-conditions shocks are proxied by CLIFS, which measures the materialisation of financial risk. An increase in CLIFS corresponds to a tightening of financial conditions.

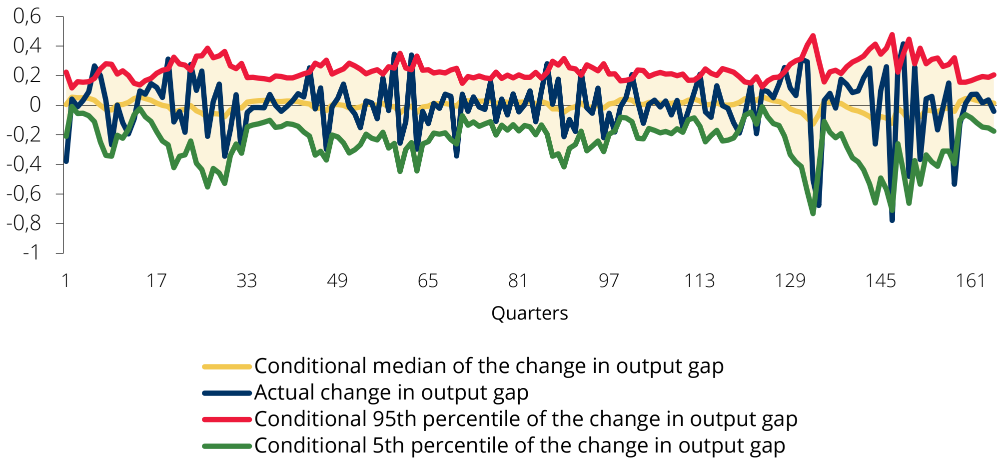

FIGURE 5: Simulation of the changes in the output gap | % deviation from the steady state

Note: the *actual* output gap change is the change computed in a benchmark simulation. Estimated percentiles are computed relative to the benchmark simulation.
Source: Banco de Portugal, authors' calculations.

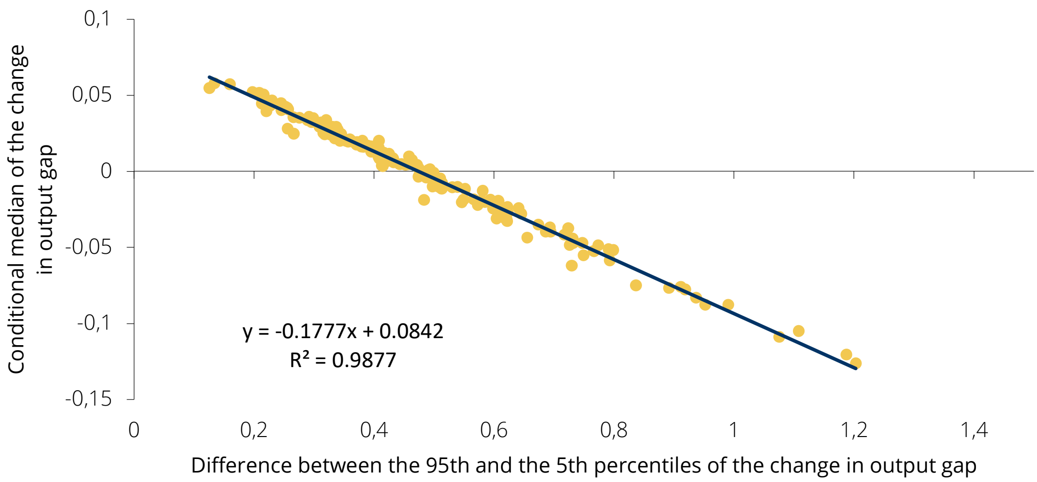

FIGURE 6: Correlation of the simulated output gap change with its volatility | % deviation from the steady state

Note: estimated percentiles are computed relative to a benchmark simulation.
Source: Banco de Portugal, authors' calculations.

with non-constant volatility and more volatile financial conditions than in the baseline specification.[7]

7. More volatile financial conditions are obtained by imposing 3-times higher shocks to the variance of financial-conditions shocks processes. Such shocks result in heavier tails in the simulated distribution of the financial conditions variable η_t.

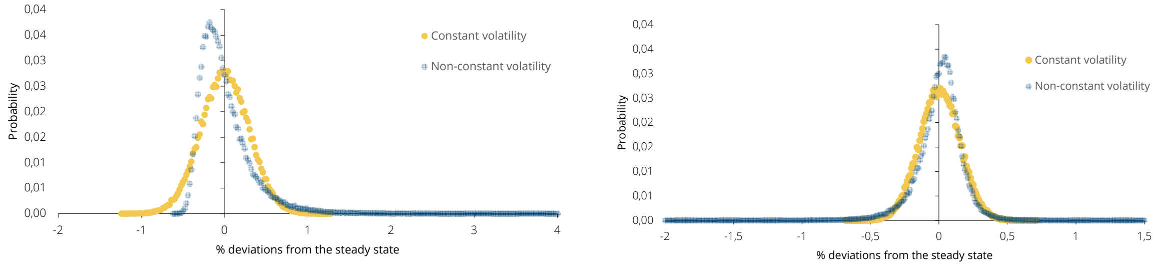

(A) Output gap (level) (B) Output gap (change)

FIGURE 7: Simulated distributions – non-constant vs constant volatility | probability
Source: Banco de Portugal, authors' calculations.

Figure 8 shows that the presence of more volatile financial conditions amplifies the intertemporal risk-taking dynamics in the model (panel in the left). In addition, the panel in the right shows that the simulated distribution of the output gap in the alternative specification with non-constant volatility and more volatile financial conditions is more volatile than in the baseline specification with non-constant volatility and less volatile financial conditions.

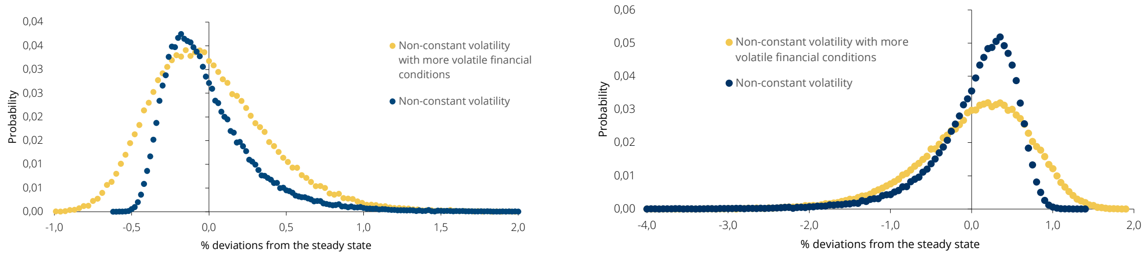

(A) Financial conditions (η_t) (B) Output gap (change)

FIGURE 8: Simulated distributions – non-constant vs constant volatility with more volatile financial conditions | probability
Source: Banco de Portugal, authors' calculations.

With more volatile financial conditions in fact, low volatility of the output gap today results in overall more volatile realizations of the output gap in the future compared to baseline specification with non-constant volatility. These dynamics translate into a more dispersed distribution of the output gap, with a higher probability of observing extreme values (positive and negative). The model in the specification with constant volatility is unable to capture this result, and for this reason, the non-constant volatility feature should be taken into account by policymakers.

5. Conclusions

Some recent macroeconomic literature is raising awareness on the need for adapting existing – or developing new – structural methodological frameworks to understand how severe macroeconomic downturns could be in the case of extreme events. The NKV model proposed in Adrian *et al.* (2020) is a pioneer approach to include a financial vulnerability channel in a structural macroeconomic model and embodies a significant starting point to study the impact of non-linearities in these types of frameworks.

In the NKV model estimated for Portugal, we replicate the non-linearities which are evident in the data, with financial conditions affecting lower realizations more than higher realizations of the output gap change, thus driving its volatility. In addition, low realizations of the output gap change are associated with periods in which its extreme negative realizations are more likely to occur. The model highlights a crucial intertemporal risk trade-off faced by policymakers, as low volatility of the output gap today triggers excessive risk-taking and results in overall more volatile realizations of the output gap in the future. Our results show that reducing the volatility of financial conditions can alleviate the intertemporal risk-taking dynamics.

Although the NKV model is limited to a short-term perspective since it is specifically designed to match one-quarter-ahead dynamics, it constitutes an important added value for policymakers, as reproduces agents' risk-taking intertemporal behaviour and the interaction of key macroeconomic variables, while including inflation and monetary policy. Therefore, this framework is useful to assess the financial stability implications of accommodative monetary policy, but also monetary policy rules that take into account financial stability risks throughout the financial cycle, by tightening to *lean against the wind* when cyclical systemic risk builds up and loosening to "clean" when systemic risk materialises. In turn, the model can also be extended to integrate macroprudential policy, which, alongside microprudential policy, is defined as the first line of defense against the build up of financial stability risks. Macroprudential measures may be better suited to deal with intertemporal risk-taking trade-offs generated by accommodative monetary policy stance, because it is specifically designed to address the financial markets imperfections and externalities which cause them, complementing monetary policy actions, which is particularly relevant in a monetary union. Countercyclical macroprudential policy rules, such as the ones mimicking countercyclical capital buffers, could be assessed against the worsening of financial imbalances leading to higher capital buffers and resilience of the institutions. In this context, the effectiveness of these policy rules in reducing the medium-term financial stability risks created by monetary policy accommodation can be examined. Finally, the interactions between the two policy domains and their role in stabilising financial conditions can also be explored, complementing growth-at-risk empirical analysis to support policy decisions.

References

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone (2019). "Vulnerable Growth." *American Economic Review*, 109(4), 1263–1289.

Adrian, Tobias, Fernando Duarte, Nellie Liang, and Pawel Zabczyk (2020). "NKV: A New Keynesian Model with Vulnerability." *AEA Papers and Proceedings*, 110, 470–476.

Aikman, David, Kristina Bluwstein, and Sudipto Karmakar (2023). "A tail of three occasionally-binding constraints: a modelling approach to GDP-at-Risk." Working paper 931, Bank of England.

Aikman, David, Jonathan Bridges, Sinem Hacioglu Hoke, Cian ONeill, and Akash Raja (2019). "Credit, capital and crises: a GDP-at-Risk approach." Working paper 824, Bank of England.

An, Sungbae and Frank Schorfheide (2007). "Bayesian analysis of DSGE models." *Econometric Reviews*, 26(2-4), 113–172.

Brownlees, Christian and André B.M. Souza (2021). "Backtesting global Growth-at-Risk." *Journal of Monetary Economics*, 118, 312–330.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2024). "Capturing Macro-Economic Tail Risks with Bayesian Vector Autoregressions." *Journal of Money, Credit and Banking*, 56(5), 1099–1127.

Cesa-Bianchi, Ambrogio and Emilio Fernandez-Corugedo (2018). "Uncertainty, Financial Frictions, and Nominal Rigidities: A Quantitative Investigation." *Journal of Money, Credit and Banking*, 50, 603–636.

Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno (2014). "Risk Shocks." *American Economic Review*, 104(1), 27–65.

De Lorenzo Buratta, Ivan, Marina Feliciano, and Duarte Maia (2022a). "How Bad Can Financial Crises Be? A GDP Tail Risk Assessment for Portugal." Working Paper Series 04, Banco de Portugal.

De Lorenzo Buratta, Ivan, Marina Feliciano, and Duarte Maia (2022b). "Mind the Build-up: Quantifying Tail Risks for Credit Growth in Portugal." Working Paper Series 07, Banco de Portugal.

Duprey, Thibaut, Benjamin Klaus, and Tuomas Peltonen (2015). "Dating systemic financial stress episodes in the EU countries." Working Paper Series 1873, European Central Bank.

Fernández-Villaverde, Jesus and Juan F. Rubio-Ramírez (2007). "Estimating macroeconomic models: A likelihood approach." *Review of Economic Studies*, 74(4), 1059–1087.

Galí, Jordi (2008). *Introduction to Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

Lang, Jan Hannes, Cosimo Izzo, Stephan Fahr, and Josef Ruzicka (2019). "Anticipating the bust: a new cyclical systemic risk indicator to assess the likelihood and severity of financial crises." Occasional Paper Series 219, European Central Bank.

Montes-Galdón, Carlos, Viktors Ajevskis, František Brázdik, Pablo Garcia, William Gatt, Diana Lima, Kostas Mavromatis, Eva Ortega, Niki Papadopoulou, Ivan De Lorenzo Buratta, and Benedikt Kolb (2024). "Using structural models to understand macroeconomic tail risks." Occasional Paper Series 357, European Central

Bank.

Papadamou, Stephanos, Moise Sidiropoulos, and Aristea Vidra (2018). "A Taylor Rule for EU members. Does one rule fit to all EU member needs?" *The Journal of Economic Asymmetries*, 18, 1–1.

Passinhas, Joana and Ana Pereira (2023). "A macroprudential look into the risk-return framework of banks' profitability." Working Paper Series 03, Banco de Portugal.

Smets, Frank and Raf Wouters (2003). "An estimated dynamic stochastic general equilibrium model of the euro area." *Journal of the European Economic Association*, 1(5), 1123–1175.

Smets, Frank and Raf Wouters (2007). "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American Economic Review*, 97(3), 586–606.

Appendix:

Parameters	Description	Value	s.d.	t-stat
Linear model estimation				
δ	Parameter of the Phillips curve*	0.0952	155.3301	0.0006
λ^h	Consumption habit	0.9251	5.3964	0.1714
ρ^i	Persistence of the monetary policy shock	0.4878	295.0003	0.0017
ρ^z	Persistence of the productivity shock	0.8451	151.6026	0.0056
ρ^θ	Persistence of the cost-push shock	0.8514	2013.4155	0.0004
σ^i	Std. Dev. of the monetary policy shock	0.0047	3.6468	0.0013
σ^z	Std. Dev. of the productivity shock	0.0047	1.5143	0.0031
σ^θ	Std. Dev. of the cost-push shock	0.0050	55.3017	0.0001
Non-linear model estimation				
v_c	Parameters of the vulnerability channel equation (3)	0.1560	0.0000	9255.8116
ϱ_{η_1}	–	0.0200	0.0021	9.4477
ϱ_{η_2}	–	0.0200	0.0012	16.5690
ϱ_{gap}	–	0.3000	0.0058	51.7480

* $\kappa = \frac{(1-\theta)(1-\beta\theta)}{\theta} \cdot \frac{1}{1+\beta\delta}$

TABLE A.1. Estimated Parameters

Non-technical summary

October 2025

Free riding in climate mitigation

Bernardino Adão

When any country emits greenhouse gases, it adds to the overall level of carbon in the atmosphere, which affects the climate everywhere—not just in that one country. So, keeping world climate emissions low is something that benefits everyone, no matter who pays for it. That makes it a global public good: you can't stop anyone from enjoying it, and one country benefiting doesn't take away from others.

But that's where the problem starts. Since cutting emissions is costly—whether through stricter rules or investing in cleaner technologies—countries might try to do less and hope others pick up the slack. This is called free riding. If too many countries do this, the total effort to reduce emissions may be too weak to prevent severe climate consequences.

It can get even trickier because not all countries are affected by climate change in the same way. Some might not see big impacts, or could even benefit a little, especially those farther from the equator. Others—often those near the equator—could face major damages. This uneven impact messes with incentives and makes global cooperation harder. On top of that, international agreements don't have strong ways to enforce climate promises, unlike domestic policies.

This article explores how free-riding incentives are shaped by a country's size and the extent of regional climate damages it faces. It establishes, in the context of a simple model, various qualitative results. First, the cooperative or first-best outcome is achieved when all countries impose the same carbon tax (a Pigouvian tax) that fully internalizes the global externality caused by carbon emissions. This harmonized tax offsets the environmental harm from emissions. Second, in the non-cooperative scenario, each country sets its own carbon tax, which typically falls short of the globally optimal level. Third, holding other factors constant, smaller countries have weaker incentives to reduce emissions compared to larger ones, as their individual impact on global outcomes is more limited. Fourth, all else equal, countries that suffer more severe consequences from climate change tend to adopt higher carbon taxes, reflecting their greater stake in mitigation.

Utilizing the model and calibration framework developed by Hassler, Krusell, Olovsson and Reiter (2020), the article shows that larger economies—such as China—and those more severely affected by climate impacts—such as India—have stronger incentives to implement higher carbon taxes. Importantly and surprisingly, the findings suggest that when China enacts its optimal carbon tax while the rest of the world applies a uniform global Pigouvian tax, the resulting policy combination leads to relatively low aggregate economic costs and only a modest rise in projected global temperatures.

Figure 1 presents the annual percent difference between the annual consumption under the scenario where all regions in the world implement the global Pigouvian tax and the annual consumption under the scenario that China implements its best climate policy when all other regions implement the global Pigouvian tax. According to the figure, the economic costs associated with China opting not to implement the global Pigouvian tax appear to be relatively limited. By the close of the 21st century, the projected annual consumption loss for the European Union and the United States remains below 0.2%, and by 2200, it stays under 0.4%. Although China experiences a net gain from this decision, the benefit is modest, averaging less than 0.2% annually over the entire projection period. At the global level, aggregate economic losses are relatively small—approximately 0.03% in 2100 and 0.3% by 2200.

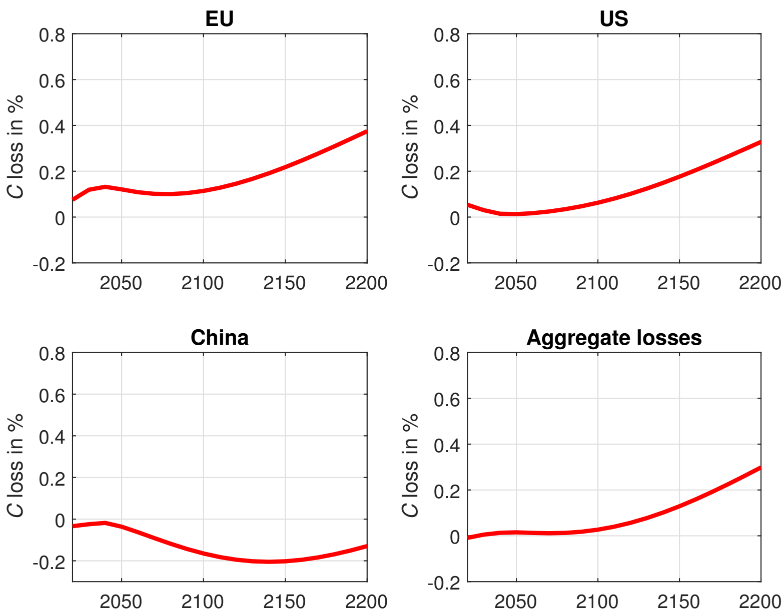

FIGURE 1: Regional welfare impacts, measured against the first-best benchmark, in the event of China's deviation.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. XI, No. 4

Free riding in climate mitigation

Bernardino Adão
Banco de Portugal

October 2025

Abstract
Carbon dioxide emissions from households and firms contribute to rising atmospheric carbon concentrations, which in turn drive global warming. Global warming represents a global externality that affects asymmetrically all regions in the world, and mitigation efforts constitute a global public good, as most regions benefit, regardless of who incurs the costs. This can create a significant free-riding problem. Some countries may opt to make only minimal reductions in their carbon emissions, choosing instead to rely on the efforts of others, as they continue to benefit from a more stable climate without assuming a proportional share of the economic responsibilities. This paper examines how free-riding incentives vary with country size and the magnitude of regional climate damages. It is established that the optimal carbon tax is independent of the strength of the climate externality in a region, and of the region's size. A common tax equalizes the marginal cost of abatement across regions, ensuring emissions are reduced where it is cheapest. In a non-cooperative equilibrium, smaller regions and regions that are less affected by climate change choose lower carbon taxes. Drawing on the model and calibration presented in Hassler *et al.* (2020), the analysis shows that larger countries, such as China, and those experiencing higher climate-related damages, such as India, have incentives to adopt carbon taxes. Notably, when China adopts its optimal carbon tax while the rest of the world implements a uniform global Pigouvian tax, the resulting policy mix yields relatively low overall economic costs and only a modest increase in global temperature projections. (JEL: E62, H23, Q43, Q54, Q58)

1. Introduction

Since the onset of industrialization, average global temperatures have increased by approximately 1°C, primarily due to carbon dioxide emissions from fossil fuel combustion.[1] This accumulation of greenhouse gases has been associated with

Acknowledgements: The author thanks the comments made by Nuno Alves, António Antunes, Ana Fontoura Gouveia and Pedro Duarte Neves. The views expressed in this paper are the author's and do not necessarily represent those of Banco de Portugal or the Eurosystem.

E-mail: badao@bportugal.pt

1. The main cause of today's climate change is the release of CO2 into the atmosphere. CO2 plays the biggest role because it has a strong and long-lasting effect on the climate. The key reason CO2 affects the climate is that it changes the Earth's energy balance – the difference between the energy the Earth receives from the sun and the energy it sends back into space. CO2 lets most sunlight pass through the atmosphere but traps heat, which leaves the Earth as infrared radiation. When there's more CO2 in the atmosphere, less heat escapes. This creates an energy imbalance, causing the planet to warm. As the Earth gets hotter,

an increased incidence of climate-related phenomena such as droughts, floods, and hurricanes. The belief among the scientific community is that the low marginal cost of extracting and using carbon-intensive energy sources implies that, absent coordinated policy interventions, emissions will likely continue along a trajectory inconsistent with long-run optimal welfare outcomes. Projections from physical climate models imply that without such policy action, the frequency and intensity of adverse weather events will rise with further warming (IPCC (2021a)).

There is a broad consensus that substantial greenhouse gas emission reductions are required. While the precise consequences of continued greenhouse gas emissions cannot be forecast with confidence, the structure of the problem does not permit us to rule out outcomes associated with large welfare losses. In the presence of such non-negligible tail risks, precautionary action becomes a rational policy response. The relevant question is whether a transition to climate neutrality can be achieved at acceptable economic cost. If so, such a transition dominates the business as usual in expected utility terms. From a policy standpoint, then, the central task is not to assign speculative values to the damages of inaction, but to rigorously characterize feasible transition paths – accounting for constraints on technology, preferences, and institutions – and to identify the incentives required to make those paths implementable in equilibrium.

Emissions of greenhouse gases by any individual country contribute to the global atmospheric stock of carbon, generating externalities that affect temperature and climate outcomes in all countries. The resulting stability of the climate thus constitutes a global public good – non-excludable and non-rival – implying that the benefits of abatement are shared irrespective of who incurs the cost. In this setting, the standard problem of free riding emerges: since mitigation entails real economic costs – whether through regulatory constraints or capital investment in low-carbon technologies – countries may have an incentive to under-provide effort, relying instead on the actions of others. If this behavior is widespread, aggregate mitigation will fall short of what is required to avoid large-scale climatic disruptions.

The difficulty is compounded by cross-country heterogeneity: the economic impact of climate change varies across countries, with some – particularly at higher latitudes – potentially experiencing limited or even positive effects, while others – often in equatorial zones – bear disproportionate losses. These asymmetric payoffs distort incentives and further complicate the implementation of coordinated global policy.

Moreover, unlike the policies designed to take care of domestic externalities, international agreements lack enforceable mechanisms. The recent decision by the US administration in January 2025 to again withdraw from the Paris Agreement illustrates the fragility of such commitments. As such, any viable policy response must be grounded in institutional arrangements that are both incentive-compatible and robust to unilateral deviation.

---

it starts to release more heat, and eventually a new balance is reached – but at a higher temperature than before.

General equilibrium theory makes clear that in the presence of externalities, decentralized market outcomes are not efficient. In the context of climate change, absent a mechanism to internalize the social cost of carbon emissions, private agents will emit beyond the socially optimal level. The canonical solution involves pricing the externality – via carbon taxes or tradable emissions permits – so as to align private incentives with social welfare.

These instruments, however, must be implemented within a dynamic framework that accounts for intertemporal trade-offs, substitution across inputs, and endogenous technological progress. Yet the effectiveness of such mechanisms at the global scale depends critically on institutional enforcement. In the absence of binding and credible international agreements, individual countries face incentives to deviate, especially when large emitters – such as China, India, or the United States under recent leadership – choose not to participate or to scale back commitments. Under these conditions, even well-designed policy instruments may fail to achieve globally efficient outcomes.

In this paper we look at non-cooperative outcomes for climate change. In this scenario some countries act in their own self-interest without coordinating efforts to reduce global greenhouse gas emissions. These outcomes typically result from a failure to collaborate or agree on shared responsibilities, and they can significantly worsen the overall climate crisis.

We start by establishing, in the context of a simple model, various qualitative results. First, the cooperative solution, or the first-best solution, is obtained with a carbon tax (or Pigou tax) equal across countries. This tax cancels the negative externalities from the emissions that carbon use causes. Second, in the non-cooperative solution the carbon tax chosen by each country is below the first-best level. Third, ceteris paribus, smaller countries have a lower incentive to cut back carbon emissions than large countries. Fourth, ceteris paribus countries that are more adversely affected by climate change choose higher carbon taxes.

Later, in the context of the Hassler *et al.* (2021b) dynamic quantitative model, which has 8 regions, we establish that large countries, like China, and countries with relatively large climate damages, like India, choose a carbon tax that is a non-negligible fraction of the optimal carbon tax. The best response carbon tax of these countries is larger the lower is the subjective discount factor of the consumers preferences and the lower is the carbon tax of the rest of the world.

Hassler *et al.* (2024) finds that "compensating for large regions that fail to introduce policies limiting fossil fuel use is very costly, or downright impossible," and that "successful climate policy requires that all regions of the world participate." Our analysis suggests that regions have an incentive to participate. For example, if China deviates from the uniform global Pigouvian tax and instead adopts its own optimal climate policy, while all other regions adhere to the global tax, the resulting aggregate costs are low and the global temperature trajectory is largely unaffected.

This outcome is explained by the nature of the private marginal cost of taxation, which increases with the tax rate – it is zero at a zero tax rate and grows progressively steeper. As shown in Hassler *et al.* (2021b) even a relatively low global carbon tax can lead to substantial emission reductions, highlighting that full participation at high tax levels may not be strictly necessary for effective climate mitigation.

The remainder of the paper is organized as follows. Section 2 considers a static framework in which the underlying intuitions are more easily grasped. Section 3 presents the dynamic, data-calibrated model that supports the analysis. Section 4 presents the results for the different scenarios under consideration. Section 5 concludes.

2. A static model

For clarity, we consider a simplified model to develop intuition before moving to the general case. The static economy considered here is a simple version of the one studied in Section 3. There are N regions and one representative consumer in each region, whose welfare is determined by consumption. Each region has size ω_i, where $\omega_i > 0$ and $\sum_{i=1}^{N} \omega_i = 1$. The production Y_i of the consumption good is given by $Y_i = F(K_i, L_i, E_i)$, where the inputs K_i, L_i , and E_i are capital, labor and energy used in region i. Function F has constant returns to scale. In the static model K_i, L_i are not endogenous and are normalized to ω_i. Thus, $Y_i = F(\omega_i, \omega_i, E_i) = \omega_i F\left(1, 1, \frac{E_i}{\omega_i}\right) \equiv \omega_i f(\tilde{E}_i)$, where $\tilde{E}_i \equiv \frac{E_i}{\omega_i}$, is the per household consumption of energy in region i. Energy is produced from two inputs: a fossil-based, or "dirty", input e_1 and another, non-fossil, or "clean", input e_2. We assume

$$E_i = G(e_1^i, e_2^i), \tag{1}$$

where G is a constant elasticity of substitution (CES) and constant returns to scale function. We assume the production of these two inputs, e_1^i and e_2^i has constant marginal costs, p_1 and p_2, respectively, in terms of output.

Later, we follow the literature by assuming that damages from climate change lowers total-factor productivity of the final good. For now, we consider damages as simply additively decreasing consumption. The per-capita damages are given by the expression $\gamma_i D(\sum_{i=1}^{N} e_1^i)$. The γ_i is a parameter that determines the magnitude of the damage costs and $D(\sum_{i=1}^{N} e_1^i)$ is a convex, twice differentiable function of the world consumption of the "dirty" input, $\sum_{i=1}^{N} e_1^i$. Thus, $\gamma_i D(\cdot)$ summarizes how the use of the aggregate "dirty" input causes emissions, which affect the carbon concentration in the atmosphere, leading to harmful global warming in region i. Damage costs in any one region depend on the aggregate level of emissions across all regions (emissions are a global public bad). We allow the damages to be different across regions because changes in the atmospheric carbon concentration affect each region's climate differently. Region i resource constraint is

$$\omega_i C_i + p_1 e_1^i + p_2 e_2^i = \omega_i f(\tilde{E}_i) - \omega_i \gamma_i D\left(\sum_{i=1}^{N} e_1^i\right), \tag{2}$$

where C_i is the consumption of the representative household in region i.

We allow for a carbon tax. It is the instrument that is usually considered in the literature and when set at the correct level it eliminates the negative externality associated with the carbon emissions.[2]

There are three levels in the production chain and a representative firm in each of these levels. The problem of the representative firm that produces the final good is to choose the input E_i that maximizes profits:

$$\omega_i f(\tilde{E}_i) - P_i E_i, \tag{3}$$

where P_i is the price of energy in country i. The problem of the representative firm that produces energy is to choose the vector $\{E_i, e_1^i, e_2^i\}$ that maximizes

$$P_i E_i - (p_1 + \tau_i) e_1^i - p_2 e_2^i, \tag{4}$$

subject to

$$E_i = G(e_1^i, e_2^i), \tag{5}$$

where τ_i is the carbon tax in region i on the use of the "dirty" input. Finally, as referred above "dirty" and "clean" sources of energy are produced at constant marginal costs (in terms of the final output good), p_1 and p_2, by a third kind of firm and sold at production cost to the firms that produce energy. The revenues from taxes are handed to consumers, so that aggregate consumption in region i is

$$\omega_i C_i = \omega_i f(\tilde{E}_i) - P_i E_i + \tau_i e_1^i. \tag{6}$$

The competitive equilibrium in region i is a vector of allocations, (C_i, E_i, e_1^i, e_2^i), prices (P_i, p_1, p_2) and carbon tax, τ_i, that satisfy a set of conditions. Given the prices and the tax, (P_i, p_1, p_2, τ_i), the allocations (E_i, e_1^i, e_2^i) solve the problems of the firms, and C_i satisfies the resource constraint (2).

Now we characterize the competitive equilibrium. Given that the production function of energy is homogeneous of degree one, the costs are linear in the level of energy, and the price of energy is $P_i = \mathbf{C}(\tilde{p}_{1,i}, p_2)$, where $\tilde{p}_{1,i} \equiv p_1 + \tau_i$ and $\mathbf{C}$ is the unit cost function.[3] From Shepard's lemma, the demand functions for the inputs are

$$e_1^i = \mathbf{C}_1(\tilde{p}_{1,i}, p_2) E_i, \tag{7}$$

and

$$e_2^i = \mathbf{C}_2(\tilde{p}_{1,i}, p_2) E_i. \tag{8}$$

2. Other instruments are possible. A popular one is a regulatory system where only a certain amount of carbon emissions is allowed, with well defined property rights to emit that can be traded. In the model, such a "cap and trade" system would be equivalent to a carbon tax system: for any carbon tax there is a maximum emissions level yielding the same equilibrium allocation, and vice versa.

3. The unit cost function is given by $\mathbf{C}(\tilde{p}_{1,i}, p_2) = \min_{e_1^i, e_2^i} \{\tilde{p}_{1,i} e_1^i + p_2 e_2^i \text{ s.t. } G(e_1^i, e_2^i) = 1\}$.

Functions $\mathbf{C}_1$ and $\mathbf{C}_2$ are the cost-minimizing input levels that produce one unit of energy. The final-good producer chooses the E_i that maximizes profits:

$$\omega_i f(\tilde{E}_i) - \mathbf{C}(\tilde{p}_{1,i}, p_2)E_i. \tag{9}$$

The first order condition of this problem is:

$$f'(\tilde{E}_i) = \mathbf{C}(\tilde{p}_{1,i}, p_2). \tag{10}$$

Equation (10), together with the demand for the inputs, (8) and (7) uniquely define E_i, along with e_1^i and e_2^i.

Now, we assume that there is a single social planner choosing the optimal levels of mitigation for the world. This assumption would be equivalent to assuming a world government, or that all regions work together to address climate change, for example through a binding global international environmental agreement. We suppose that the utility function of the representative household of region i is a function of his or her consumption, and the utility functions across countries are equal and linear in consumption, $u(C_i) = C_i$. The consumption of the representative household of country i is

$$C_i = f(\tilde{E}_i) - p_1 \frac{e_1^i}{\omega_i} - p_2 \frac{e_2^i}{\omega_i} - \gamma_i D\left(\sum_{i=1}^{N} e_1^i\right). \tag{11}$$

We define the cooperative solution as the vector of allocations $\left(E_i, \tau_i, e_1^i, e_2^i\right)_{i=1}^{N}$ that solves the following problem:

$$\max \sum_{i=1}^{N} \omega_i C_i \tag{12}$$

subject to

$$f'(\tilde{E}_i) = \mathbf{C}(\tilde{p}_{1,i}, p_2), \text{for } i = 1, ..., N, \tag{13}$$

$$e_1^i = \mathbf{C}_1(\tilde{p}_{1,i}, p_2)E_i, \text{for } i = 1, ..., N, \text{and} \tag{14}$$

$$e_2^i = \mathbf{C}_2(\tilde{p}_{1,i}, p_2)E_i, \text{for } i = 1, ..., N. \tag{15}$$

As shown in Appendix A, the carbon tax that solves this problem is:

$$\tau^{CO} = D'\left(\sum_{j=1}^{N} e_1^{j,CO}\right) \sum_{i=1}^{N} \omega_i \gamma_i, \tag{16}$$

where $e_1^{j,CO}$ is the "dirty" input used in region j when all regions implement the first-best carbon tax. This equation delivers our first result.

Result 1: The carbon tax in the cooperative solution is equal in all regions. It is independent of the size of the region and its damage parameter.

The rationale for a uniform carbon tax across countries is justified by considerations of economic efficiency and the global nature of climate change. A common tax equalizes the marginal cost of abatement (i.e., the cost of reducing one additional ton of emissions) across regions, ensuring emissions are reduced where it is cheapest. Since CO2 has the same climate impact regardless of origin, equal pricing makes environmental sense. Uniform taxation also helps prevent carbon leakage – where firms shift operations to low-tax regions – and is simpler to implement and enforce than a fragmented policy approach.

Next, we characterize the carbon tax that maximizes the welfare of region i, when there are no external rules or binding agreements to enforce cooperation between regions. As before, the tax is paid by firms and the payments are rebated to the consumers. Taking as given the choices in the other regions, $\left(E_j, \tau_j, e_1^j, e_2^j\right)_{j \neq i}$, the optimal level of the tax in region i must be a solution to the following problem:

$$\max \ \omega_i C_i \tag{17}$$

subject to

$$f'(E_i) = \mathbf{C}(\tilde{p}_{1,i}, p_2), \tag{18}$$

$$e_1^i = \mathbf{C}_1(\tilde{p}_{1,i}, p_2) E_i, \text{ and} \tag{19}$$

$$e_2^i = \mathbf{C}_2(\tilde{p}_{1,i}, p_2) E_i. \tag{20}$$

As we show in Appendix A, the carbon tax that solves this problem is

$$\tau^{i,NC} = D' \left(\sum_{j=1}^{N} e_1^{j,NC} \right) \omega_i \gamma_i. \tag{21}$$

where $e_1^{j,NC}$ is the "dirty" input used in region j when region j implements the carbon tax $\tau^{j,NC}$, for $j = 1, ..., N$. This equation delivers our second result.

Result 2: In the non-cooperative equilibrium, smaller regions (i.e., regions with smaller ω_i) and regions that are less affected by climate change (i.e., regions with smaller γ_i) choose lower carbon taxes.

The intuition of this result is clear. Smaller regions contribute relatively little to total global emissions, so their individual efforts to cut emissions have little impact on global temperatures. However, they still bear the full cost of any climate policies they implement. At the same time, they benefit from the emission reductions of larger regions, since climate stabilization is a global public good. This creates a strong incentive for small regions to free-ride – avoiding the costs of mitigation while still enjoying its benefits.

Regions more affected by climate change have more to lose from CO2 emissions. For them, the benefits of cutting emissions, in the form of avoided damages, outweigh the costs of climate policies. As a result, even if other regions do not act, these more vulnerable regions have a stronger incentive to reduce their own emissions. Similarly, regions that are less adversely affected by climate change will prefer to emit more.

In order to further compare the cooperative and non-cooperative carbon taxes, it is convenient to assume specific functional forms. That is what we do now. Assume the production of the final good is:

$$f(E) = B(E)^{\varphi}, B > 0, \text{and } 0 < \varphi < 1, \tag{22}$$

where φ is the share of energy in the production of the final good, and that the production of energy is:

$$E = G(e_1, e_2) = e_1^{\alpha} e_2^{1-\alpha}, 0 < \alpha < 1, \tag{23}$$

where α is the share of the "dirty" source of energy. The cost function is given by

$$\mathbf{C}(\tilde{p}_1, p_2, E) = \min_{e_1, e_2} \{\tilde{p}_1 e_1 + p_2 e_2 \text{ s.t. } e_1^{\alpha} e_2^{1-\alpha} = E\}, \tag{24}$$

The damage function is a quadratic function:

$$D(x) = \left(\frac{x^2}{2}\right), \text{ where } x = \sum_{j=1}^{N} e_1^j. \tag{25}$$

As we show in Appendix A, for these functional forms, the first-best carbon tax is such that:

$$\tau^{CO} = \sum_{j=1}^{N} \omega_j \gamma_j \Phi p_2^{-\frac{\varphi(1-\alpha)}{1-\varphi}} \sum_{k=1}^{N} (\tilde{p}_1^{CO})^{-\frac{1-\varphi(1-\alpha)}{1-\varphi}} \omega_k, \tag{26}$$

where $\Phi = \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} \left(\frac{A}{B\varphi}\right)^{\frac{1}{\varphi-1}}$ and $A = \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} + \left(\frac{\alpha}{1-\alpha}\right)^{-\alpha}$. This equation provides additional insights that we summarize in Result 3.

Result 3: The optimal carbon tax is higher the lower are the prices of the inputs p_1 and p_2.

Result 3 is intuitive. The lower the input prices, p_1 and p_2, the lower the overall cost of energy production. Consequently, a reduction in energy costs encourages greater use of energy in the production of the final good. To counteract this increased energy consumption and effectively curb greenhouse gas emissions, the carbon tax must be correspondingly higher.

As we show in Appendix A, for the functional forms adopted, the non-cooperative carbon taxes satisfy the equations:

$$\tau^{i,NC} = \omega_i \gamma_i \Phi p_2^{-\frac{\varphi(1-\alpha)}{1-\varphi}} \sum_{k=1}^{N} (\tilde{p}_1^{k,NC})^{-\frac{1-\varphi(1-\alpha)}{1-\varphi}} \omega_k, \text{for } i = 1, ..., N. \tag{27}$$

The direct comparison between the optimal cooperative tax and the noncooperative carbon tax, based on the two equations, (26) and (27), implies that $\tau^{i,NC} < \tau^{CO}$. The intuition for this result follows from the fact that if $\tilde{p}_1^{k,NC} = \tilde{p}_1^{CO}$ then the right-hand side of (27) would be smaller than the right hand side of (26). We summarize this finding as result 4.

Result 4: The non-cooperative carbon tax is smaller than the optimal carbon tax.

The intuition is that, in a non-cooperative setting, each region sets its carbon tax considering only its own costs and benefits. Since the benefits of reducing emissions are shared worldwide, an individual region captures only a fraction of the benefits from its own abatement efforts. This leads to underinvestment in emissions reduction, as regions prefer to "free-ride" on the efforts of others.

3. Dynamic model

3.1. Economy

Here, we follow the influential model developed by Golosov *et al.* (2014) and used in, for instance, Hassler *et al.* (2016), Krusell *et al.* (2020), Hassler *et al.* (2020), Hassler *et al.* (2021b), Hassler *et al.* (2024), and Adão *et al.* (2024). It presents a unified macroeconomic framework to understand how economic activity, emissions, and climate change interact over time. The model combines economic growth with climate damages, aiming to highlight the optimal climate policy in a general equilibrium context. The model includes infinitely lived households that choose consumption and investment to maximize lifetime utility. Their preferences are over consumption with standard time discounting. Output is produced using capital, labor and energy, following a neoclassical production function. Emissions of carbon are proportional to the amount of fossil fuels used to produce energy. The climate externality arises from the emissions that accumulate in the atmosphere and raise the global temperature. The temperature increases cause economic damages, modeled as a reduction in output (or total factor productivity). The model captures the interaction between emissions, temperature rise, and economic output, and shows how optimal policy can align private incentives with social welfare.

The model contains eight regions. One of the eight regions produces conventional oil, which it sells at a competitive global world market. This region has no other production and is the only producer of conventional oil. The price of conventional oil reflects its scarcity only: conventional oil is in limited supply and is costless to extract. This region is calibrated to represent the OPEC members and Russia. The remaining seven regions represent Europe, the United States, China, South America, India, Africa, and Oceania. South America includes all countries in South America except the countries that are part of OPEC, and the same is true for Africa. The region referred to as India also includes Pakistan and Bangladesh. The definition of Oceania, finally, is also wider and somewhat different from the geographical region: it includes Australia, Japan, Indonesia, Malaysia, Myanmar, New Zealand, Philippines, Thailand, and Vietnam.

These other regions have no endowment of conventional oil. Instead, they import oil from the oil producer region, and pay for this import with a common final good that is identical in all regions. These oil consuming regions are heterogeneous in size, productivity, and initial capital stocks, though they share several structural features. There are no international capital markets. This abstraction simplifies the analysis, but it is unlikely to be of primary significance. Under the calibration – where marginal products of capital are equalized across regions at time zero and preferences are assumed identical – the long-run returns to capital will also be equalized. Thus, capital flows are unnecessary in the steady state. During the transition, some divergence in marginal products will occur, but these differences are quantitatively small and do not materially affect the model's dynamics.

The preferences of the representative household in each region are given by

$$\sum_{t=0}^{\infty} \beta^t \log C_{i,t}, \tag{28}$$

where $C_{i,t}$ is the consumption of the representative household of region i, in period t, with $\beta \in (0,1)$. The assumption of logarithmic consumption preferences simplifies the computations.

3.1.1. *Oil producers region*

The oil producers region is region 1. The problem for the representative oil producer is to choose how much oil to keep in the ground for next period, $t+1$, while taking the world market price of oil as given. The extraction of oil uses oil as well.[4] The total stock of oil in the ground in period t is R_t. The law of motion for the stock of oil is given by

$$R_{t+1} = R_t - (1+\zeta) \sum_{i=2}^{8} e_{1,i,t}, \tag{29}$$

with $\zeta > 0$, where $(1+\zeta)\sum_{i=2}^{8} e_{1,i,t}$ is the oil extraction in period t, and $\sum_{i=2}^{8} e_{1,i,t}$ the oil exported, with $R_t > 0$ for all t. The total revenue of oil producers is $p_{1,t}\sum_{i=2}^{8} e_{1,i,t}$, where $p_{1,t}$ is the price of oil at date t. The budget constraint is

$$\omega_1 C_{1,t} = p_{1,t} \sum_{i=2}^{8} e_{1,i,t}, \tag{30}$$

where ω_1 is the size of region 1.

3.1.2. *Oil consuming regions*

Each oil-consuming region has an aggregate production function for the final good, $Y_{i,t}$, given by

$$Y_{i,t} = A_{i,t} L_{i,t}^{1-\alpha-\nu} K_{i,t}^{\alpha} E_{i,t}^{\nu}, \tag{31}$$

where $A_{i,t}$ is total factor productivity (TFP), $L_{i,t}$ is labor used in final good production, $K_{i,t}$ is the capital stock, and $E_{i,t}$ denotes energy services. Parameters α and ν denote

4. This assumption is necessary in order to match the carbon emissions of this region.

the capital and energy shares of capital and energy in production. The total factor productivity is taken as exogenous.

There are two empirical regularities on the use of energy. First, energy is highly complementary to capital and labor in the short run, and even well approximated by a Leontief technology in a capital-labor composite and energy (see Hassler *et al.* (2021a)). Second, the income share of energy is roughly constant in the long run. This justifies the use of a Cobb-Douglas production function in this framework as the focus here is on the long run.

Energy services are produced competitively with firms using a nested constant elasticity of substitution production function in different energy sources. Only one region, the United States, has access to nonconventional oil reserves (fracking), $e_{4,US,t}$, which is combined with conventional oil, $e_{1,US,t}$, to produce an oil composite:

$$O_{US,t} = \left(\lambda_1^{oil} e_{1,US,t}^{\rho_t^h} + \lambda_4^{oil} e_{4,US,t}^{\rho_t^h}\right)^{1/\rho_t^h}, \tag{32}$$

where $\lambda_1^{oil} + \lambda_n^{oil} = 1$ and $(1-\rho_t^h)^{-1}$ is the elasticity of substitution between the two inputs. The other regions do not engage in fracking and $O_{i,t}$ is just given by $e_{1,i,t}$.

This nest, according to Hassler *et al.* (2021b), has a high elasticity of substitution, 10, within it – fracked and conventional oil are very close substitutes. In a second nest, the oil composite is combined with coal and green energy to produce energy services with a lower elasticity of substitution, 0.95, according to Hassler *et al.* (2021b).[5] The supply of conventional oil comes from the oil-producing region. An additive, region-specific carbon tax is applied to the use of conventional oil, fracked oil, and coal. Except for oil, all energy sources are produced with a production technology that is linear in the final good. It is assumed that $p_{k,i,t}$ units of the final good are necessary to produce one unit of energy source of type k in region i and period t.

The model incorporates a unified world market for oil, with trade balance imposed across regions. This effectively rules out intertemporal trade, an abstraction that simplifies the analysis without loss of generality. Trade in coal and green energy is permitted but plays no allocative role beyond equalizing production costs across regions. As such, these energy markets serve primarily to harmonize marginal costs, without affecting the model's central dynamics.

The energy services are produced using energy inputs, $e_{k,i,t}$, according to a uniform elasticity production function:

$$E_{i,t} = \left(\lambda_1 O_{i,t}^{\rho_t} + \sum_{k=2}^{3} \lambda_k e_{k,i,t}^{\rho_t}\right)^{1/\rho_t}, \tag{33}$$

where $O_{i,t} = e_{1,i,t}$ for $i \neq US$, $k=2$ denotes coal, and $k=3$ denotes green energy sources. The elasticity of substitution between the inputs is $(1-\rho_t)^{-1}$, which as referred above is assumed to be 0.95, and their shares sum up to one, $\sum_{k=1}^{3} \lambda_k = 1$. The energy producers

5. Higher values have been proposed in the literature, such as Acemoglu *et al.* (2023), who use a constant value of 1.86, and Papageorgiou *et al.* (2017).

in the oil consuming regions, i, minimize the costs of producing energy taking the prices of the inputs as given.

The quantity of the final good that is not used for energy production can be consumed or invested. Since the time period size is taken to be 10 years in this model, the depreciation rate of capital is taken to be 100%, which simplifies the tractability of the model. The resource constraint for each region ($i = 2, ..., 8$) is

$$\omega_i C_{i,t} + K_{i,t+1} = A_{i,t} L_{i,t}^{1-\alpha-\nu} K_{i,t}^{\alpha} E_{i,t}^{\nu} - \sum_{k=1}^{4} p_{k,i,t} e_{k,i,t} \tag{34}$$

where ω_i is the size of region i and $e_{4,i,t} = 0$ for all $i \neq US$.

3.2. *The carbon cycle*

The usage of fossil energy sources generates CO2 emissions. Emissions in period t in region i, $M_{i,t}$, are given by:

$$M_{i,t} = \sum_{k=1}^{4} g_k e_{k,i,t}, \tag{35}$$

for $i = 2, ..., 8$, and

$$M_{1,t} = g_1 \zeta \sum_{i=2}^{8} e_{1,i,t}, \tag{36}$$

where g_k is the carbon content of energy input k. For purely green energy sources the carbon content is zero, $g_3 = 0$. According to (36), the extraction of oil generates emissions too.

There is a relationship between emissions of carbon and the carbon concentration in the atmosphere. Following Golosov *et al.* (2014) and Acemoglu *et al.* (2016), the excess carbon in the atmosphere is given by

$$\tilde{S}_t \equiv S_t - \overline{S} = \sum_{j=0}^{t-t^*} (1 - d_j) M_{t-j}, \qquad t > t^*, \tag{37}$$

where S_t denotes the amount of carbon in the atmosphere in period t; $\overline{S}$ is the pre-industrial amount of carbon in the atmosphere; $d_j \in [0, 1]$ measures carbon depreciation from the atmosphere j periods after it has been emitted; and $M_{t-j} = \sum_i M_{i,t-j}$ indicates the world carbon emissions in period $t - j$. Term $1 - d_j$ represents the fraction of carbon emitted j periods ago still in the atmosphere and t^* defines the start of industrialization. Out of the carbon emitted, a fraction φ_L stays in the atmosphere forever, a fraction $1 - \varphi_0$ of the remaining emissions exits into the biosphere within a decade, and the remaining part, $(1 - \varphi_L)\varphi_0$, decays at geometric rate φ. Mathematically,

$$1 - d_j = \varphi_L + (1 - \varphi_L)\varphi_0(1 - \varphi)^j . \tag{38}$$

As explained in Archer (2005), Golosov *et al.* (2014) and Acemoglu *et al.* (2016), this specification approximates the complex dynamics of carbon concentration in the atmosphere and can provide a good match to the observed dynamics of atmospheric carbon concentration; see Adão *et al.* (2022), subsection 4.1.3. More recent macroeconomic models of climate change, like Dietz *et al.* (2021) , Folini *et al.* (2025) and Fernández-Villaverde *et al.* (2025), have adopted a more complex description of the carbon circulation more in line with state-of-the-art climate science, and less with the historical data.[6]

3.3. Damages

The carbon concentration in the atmosphere determines the climate and the damages are a function of the climate. Following Hassler *et al.* (2021b), and Acemoglu *et al.* (2016), we assume that the effect of the CO2 concentration on productivity is

$$A_{i,t} = \exp(z_{i,t} - \gamma_{i,t}\tilde{S}_{t-1}), \tag{39}$$

where $z_{i,t}$ is a productivity trend, and $\gamma_{i,t}$ is a region-specific parameter that determines how climate-related damages depend on the level of the atmospheric CO2 concentration. Note that this specification implies that the marginal damage per unit of excess carbon in the atmosphere is a constant share $\gamma_{i,t}$ of net-of-damage output. Damages as a share of GDP, $D(\cdot)$, can be defined as

$$D(S_{t-1}) \equiv 1 - \exp(-\gamma_{i,t}\tilde{S}_{t-1}). \tag{40}$$

The value $\gamma_{i,t}$ is positively affected by three parameters. The first is the sensitivity of the global mean temperature to changes in the CO2 concentration, which we describe in (41) below. The second is the sensitivity of the regional climate to global mean temperature. Finally, it depends on the vulnerability of the regional economy to climate impacts.

Hassler *et al.* (2021b) adopt the simplest description of the climate, the global mean atmospheric temperature near the surface and use a relationship that relates the carbon concentration and the global mean temperature. They assume the climate system follows the energy budget model in DICE/RICE:

$$T_t - T_{t-1} = \sigma_1 \left(\frac{\eta}{\log 2} \log\left(\frac{S_{t-1}}{\overline{S}}\right) - \kappa T_{t-1} - \sigma_2(T_{t-1} - T^L_{t-1}) \right), \tag{41}$$

and

$$T^L_t - T^L_{t-1} = \sigma_3(T_{t-1} - T^L_{t-1}) \tag{42}$$

where T_t is the global mean temperature in the atmosphere and T^L_t is the mean temperature in the deep oceans. Both these temperatures are measured as deviations from their pre-industrial levels.

6. In these papers the cycle module system has three linear difference equations, each describing the change in the size of three reservoirs of carbon (often called carbon sinks). The carbon sinks are: The atmosphere, the biosphere and the surface ocean (combined), and the deep oceans. The dynamics of the reservoirs are given by a linear system where flows are proportional to the size of the source reservoirs and emissions.

The damages are calibrated using the estimates from RICE, Nordhaus and Moffat (2017). Specifically, this study presents linear-quadratic regional damage functions that quantify the share of GDP lost as a function of changes in global mean temperature. Hassler *et al.* (2021b) use the Arrhenius equation,

$$T_t = \frac{\eta}{\log 2} \log\left(\frac{S_{t-1}}{\overline{S}}\right), \tag{43}$$

for each set of damage parameters, to express the damage elasticity as a function of the global mean temperature:

$$\gamma_i = -\frac{\ln(1 - (\varphi_{1,i}T + \varphi_{2,i}T^2))}{S_0(e^{T\frac{\ln 2}{\eta}} - 1)} \quad \text{for} \quad i = 2, ..., 8. \tag{44}$$

The γ_is are obtained by taking $T = 3$°C and $S_0 = 850$ GtC in equation (44).

3.4. Governments

We assume governments implement climate policies by choosing carbon taxes, $\tau_{k,i,t}$. The objective of the study is to analyze the quantitative differences between the optimal cooperative climate policy and other non-cooperative climate policies. For simplicity, we assume the government deficit is zero and net tax revenues are transferred back to the households within the period in the form of a negative income tax rate, $\Gamma_{i,t}$, over household's income, which is given by $w_{i,t}L_{i,t} + r_{i,t}K_{i,t}$, where $w_{i,t}$ and $r_{i,t}$ denote the competitive prices of labor and capital. The government budget constraint is:

$$\Gamma_{i,t}(w_{i,t}L_{i,t} + r_{i,t}K_{i,t}) = \sum_{k=1}^{4} \tau_{k,i,t} g_k e_{k,i,t}. \tag{45}$$

3.5. Equilibrium

The households and the firms in all regions are price takers and markets operate under perfect competition. The equilibrium is defined as a set of paths for quantities and prices such that: (i) households' choices of investment and consumption maximize their utility, subject to budget constraints and taking prices as given; (ii) firms choose output and input levels that maximize profits, also taking prices as given; and (iii) all markets clear.

3.6. Optimal carbon tax

According to the intuition developed in section 2, the Pigouvian tax, or optimal carbon tax, should be equal to the costs associated with the negative externality. According to Golosov *et al.* (2014), the optimal tax per unit of carbon emitted is:

$$\tau_t = \sum_{j=0}^{\infty} \beta^j \frac{u'(C_{w,t+j})}{u'(C_{w,t})} \left(-\frac{\partial Y_{w,t+j}}{\partial S_{t+j}}\right) \frac{\partial S_{t+j}}{\partial M_t}, \tag{46}$$

where $C_{w,t}$ and $Y_{w,t}$ are the consumption of the representative household of the world economy and world production, respectively.

Given the assumptions made this expression can be simplified:

$$\tau_t = \gamma_w Y_{w,t} \left(\frac{\varphi_L}{1-\beta} + \frac{(1-\varphi_L)\varphi_0}{1-(1-\varphi)\beta} \right). \tag{47}$$

where γ_W and $Y_{w,t}$ are the world damages and the world output respectively. This equation implies that the carbon tax must grow at the world GDP growth rate. Hassler *et al.* (2021b) set the initial global GDP to 80 trillion USD and γ_W equal to 3.21×10^{-5}.

4. Results

The calibration of the parameters, as outlined in Hassler *et al.* (2021b), is described in Appendix B. In this section, the main findings from the dynamic model are reported, focusing on the contrast between first-best and non-cooperative outcomes.

4.1. Best response carbon tax

We start by defining the best response policy of a region when the remaining regions implement the global first-best policy. The best-response policy is also restricted to be a carbon tax. Let τ_i be the carbon tax implemented in region i, let τ^{opt} be the global first-best carbon tax, and let $\{C^*_{j,t+k}\}_{k=0}^{\infty}$ for $j = 1, ..., 8$, be the competitive equilibrium consumptions for the various regions, when region i implements the carbon tax τ_i and the remaining regions implement the carbon tax τ^{opt}. The consumption in any of the regions is a function of the carbon taxes, and we denote this by writing $C^*_{j,t}(\tau_i, \tau^{\text{opt}})$ for $j = 1, ..., 8$. The best response policy of region i, τ_i^{br}, is defined as:

$$\tau_i^{\text{br}} = \text{argmax}_{\tau_i} \sum_{k=0}^{\infty} \beta^{10k} \log C^*_{i,t+10k}(\tau_i, \tau^{\text{opt}}) \text{ s.t. } C^*_{i,t+10k} \text{ being an equilibrium consumption.}$$

We conduct two counterfactual experiments under the assumption that all regions except one adopt the global first-best carbon tax. In the first, region i is assigned to India; in the second, to China. In both cases, we examine the optimal unilateral response of the excluded country. The results indicate that India's optimal carbon tax is 22% of the global first-best level, while China's is 27%. This asymmetry highlights differences in regional incentives and underscores the challenge of achieving coordinated climate policy when countries act independently.[7]

4.2. Global warming under different scenarios

We now present the simulation results and associated costs for three distinct climate policy scenarios. The first is the status-quo scenario, representing the benchmark, in which no measures are taken to mitigate climate change and all regions apply zero carbon taxes. The second scenario introduces the global Pigouvian tax, which starts at \$87.10 per ton of carbon and increases annually by 2.1%, roughly in line with global GDP

7. A lower discount factor, β, leads to a lower optimal carbon tax, prompting both countries, China and India, to respond with higher carbon tax levels as their best response.

growth. This policy fully internalizes the climate externality. The third scenario explores a case in which China adopts its first-best policy while all other regions implement the global Pigouvian tax. Our objective is to assess whether this asymmetric policy compromises the effectiveness of global climate change mitigation.

Figure 1 illustrates the projected increases in global mean temperature over the next 200 years under the three policy scenarios. Several key insights emerge from the figure. First, implementing global fossil fuel taxes proves to be an effective strategy for mitigating climate change. By the end of this century, the temperature difference between the status-quo scenario and the global Pigouvian tax scenario reaches 0.65°C, with the gap widening significantly thereafter. Surprisingly, however, the temperature trajectories under the global Pigouvian tax and the third scenario – in which China adopts its first-best policy while others implement the Pigouvian tax – are nearly indistinguishable by the end of the century, and with a gap of only 0.22°C in 2200. This suggests that this policy deviation has little impact on the global climate outcome.

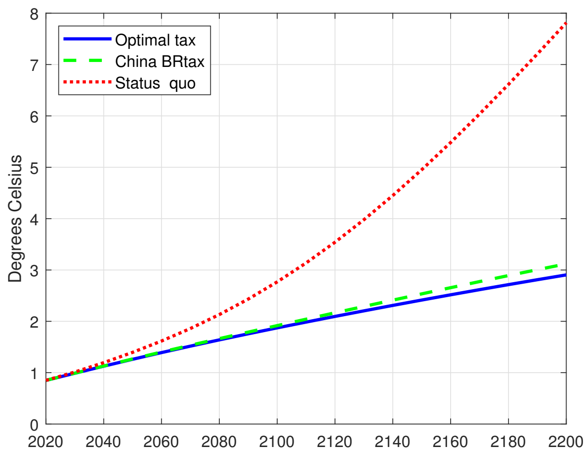

FIGURE 1: Temperature in excess of the pre-industrial baseline under various scenarios.

The reason China's deviation from the global Pigouvian tax – by adopting a lower carbon tax – has little effect on the global climate outcome is that, in this model, even modest carbon taxes lead to significant reductions in emissions. The Pigouvian tax is itself relatively modest when compared to European fuel taxes. Given gasoline's carbon content of approximately 0.65 kg per liter, the Pigouvian tax corresponds to about 5.7 cents per liter.

4.3. *Effects on consumption*

Figure 2 compares the consumptions under scenario 2 and scenario 3 for the various regions. More specifically, the figure presents the annual percent difference between annual consumption under scenario 2 and scenario 3:

$$\Delta_{i,t} = 100 \left(\frac{c_{i,t}}{c_{i,t}^{\text{opt}}} - 1 \right),$$

where $c_{i,t}$ is the annual consumption in region i in scenario 3 and $c_{i,t}^{\text{opt}}$ is the annual consumption in region i in scenario 2. The annual percent difference for the world, including the oil producers region, is

$$\Delta_t = \sum_{i=1}^{8} \Delta_{i,t} \frac{GDP_{i,t}}{GDP_{W,t}}$$

where $GDP_{i,t}$ is the GDP of region i and $GDP_{W,t}$ is global GDP.

According to the figure, the economic costs associated with China opting not to implement the global Pigouvian tax appear to be relatively limited. By the end of the 21st century, the projected annual consumption loss remains below 0.2% across all global regions. By 2200, the losses remain under 0.4% for most regions, with the exception of India and Africa, where the loss reaches approximately 0.7%. Although China experiences a net gain from this decision, the benefit is modest, averaging less than 0.2% annually over the entire projection period. Oil-exporting countries, by contrast, benefit more under Scenario 3, as their revenues increase when China adopts its first-best carbon tax response. At the global level, aggregate economic losses are relatively small – approximately 0.03% in 2100 and 0.3% by 2200. In their study, Hassler *et al.*

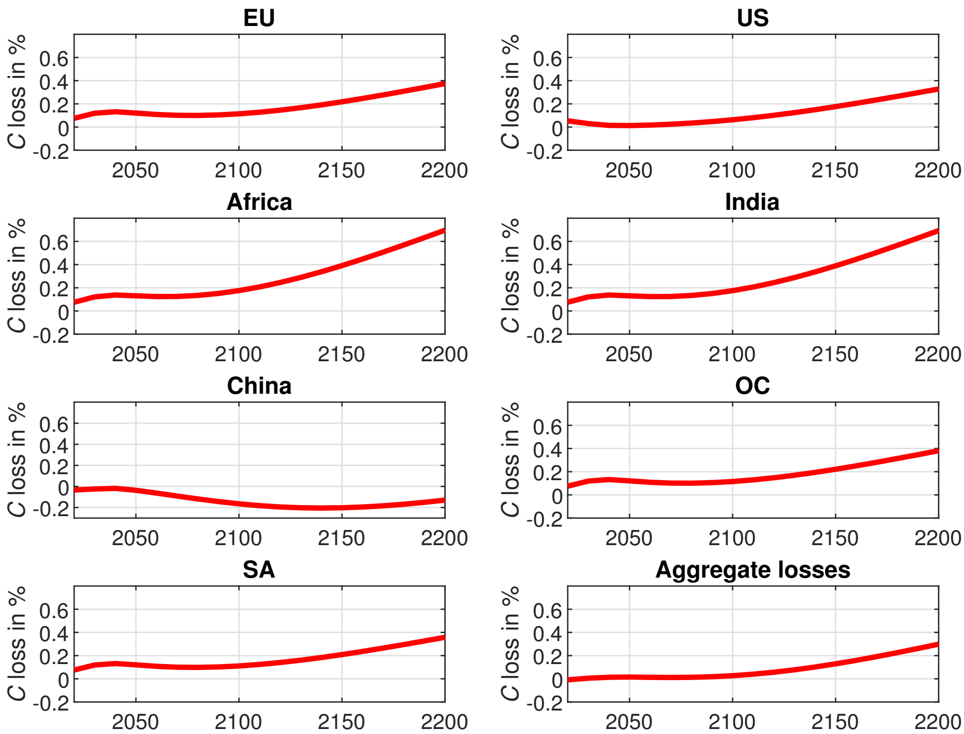

FIGURE 2: Regional welfare impacts, measured against the first-best benchmark, in the event of China's deviation.

(2020) explore a scenario where China participates in climate policy only to a limited degree – implementing just 15% of the Pigouvian carbon tax – while all other regions impose a tax 20 times higher than the optimal level. This extreme measure is designed to keep the global temperature increase in line with that of a fully coordinated global Pigouvian tax (Scenario 2). Their findings show that this kind of imbalance leads to very high economic costs, with annual global consumption losses exceeding 5%. Based on this, they conclude that effective climate policy really depends on broad international participation, and that trying to offset the lack of action from major emitters like China is extremely costly. While this conclusion is valid, it overlooks the fact that, in certain situations, a region may actually have strong incentives to take a more active role in climate mitigation. This is especially true for regions that are either large contributors to emissions or more severely impacted by climate change than the global average.

5. Concluding remarks

This paper investigates how free-riding incentives in climate policy are influenced by country size and the severity of region-specific climate damages. Utilizing the model and calibration framework developed by Hassler *et al.* (2020), the analysis demonstrates that larger economies – such as China – and those facing greater climate-related impacts – such as India – have stronger incentives to implement carbon taxation. Significantly, when China enacts its own optimal carbon tax while the rest of the world adheres to a uniform global Pigouvian tax, the resulting policy configuration produces relatively low global economic costs and only a marginal increase in projected global temperatures.

These findings should be interpreted with caution for several reasons. First, the welfare estimates presented here do not account for the direct disutility associated with living in a world increasingly affected by extreme weather events, biodiversity loss, and other non-market damages from climate change, as emphasized by IPCC (2021b) and Adão *et al.* (2022). While incorporating such direct climate-related impacts on utility – provided they remain external to the representative agent's decision-making – would not alter the allocation outcomes described in this section, it would affect the welfare metrics, particularly under more environmentally adverse scenarios. Second, if utility were explicitly affected by climate damages, the optimal carbon tax would need to be adjusted accordingly. Specifically, the tax would include an additional term reflecting the marginal disutility of climate change, leading to a higher level of Pigouvian taxation.

Several additional caveats warrant consideration. First, the model does not account for the potential link between atmospheric carbon concentrations and the frequency or severity of extreme weather events. Moreover, the relationship between carbon levels and global temperature may be non-linear and subject to discontinuities, particularly due to potential climate tipping points. Second, the analysis abstracts from non-market damages – such as conflict-induced migration, biodiversity loss, and general discomfort – which could significantly influence welfare outcomes and undermine the ability to draw clear comparisons between scenarios.

References

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hémous (2023). "Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution." Working Paper 31657, National Bureau of Economic Research.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr (2016). "Transition to clean technology." *Journal of Political Economy*, 124(1), 52–104.

Adão, Bernardino, António Antunes, Miguel Gouveia, Nuno Lourenço, and João Valle e Azevedo (2022). "Climate change and the economy: an introduction." Occasional Paper 1, Banco de Portugal.

Adão, Bernardino, Borghan Najarabad, and Ted Temzelides (2024). "Renewable technology adoption costs and economic growth." *Energy Economics*, 129.

Archer, David (2005). "Fate of fossil fuel CO2 in geologic time." *Journal of Geophysical Research: Oceans*, 110(C9).

BP (2015). "BP Statistical Review of World Energy." `http://bp.com/statisticalreview`.

Dietz, S., F. Van Der Ploeg, A. Rezai, and F. Venmans (2021). "Are Economists Getting Climate Dynamics Right and Does It Matter?" *Journal of the Association of Environmental and Resource Economists*, 8(5), 895–921.

Fernández-Villaverde, J., K. Gillingham, and S. Scheidegger (2025). "Climate Change Through the Lens of Macroeconomic Modeling." *Annual Review of Economics*, 17, 125–50.

Folini, D., F. Kübler, A. Malova, and S. Scheidegger (2025). "The Climate in Climate Economics." *Review of Economic Studies*, 92, 299–338.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014). "Optimal taxes on fossil fuel in general equilibrium." *Econometrica*, 82(1), 41–88.

Hassler, John, Per Krusell, and Conny Olovsson (2021a). "Directed Technical Change as a Response to Natural-Resource Scarcity."." *Journal of Political Economy*, 129, 3039—-3072.

Hassler, John, Per Krusell, and Conny Olovsson (2021b). "Suboptimal Climate Policy." *Journal of the European Economic Association*, 19, 2895–2928.

Hassler, John, Per Krusell, and Conny Olovsson (2024). "The macroeconomics of climate change." Working Paper 8, PIIE.

Hassler, John, Per Krusell, Conny Olovsson, and Michael Reiter (2020). "On the effectiveness of climate policies." Working Paper 53, IIES.

Hassler, John, Per Krusell, and Anthony Smith (2016). "Environmental Macroeconomics." In *Handbook of Macroeconomics*, vol. 2, edited by John Taylor and Harald Uhlig, pp. 1893–2008. Elsevier Press.

IPCC (2021a). "Climate change 2021: the physical science basis. Summary for policymakers." In *Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by V. Masson-Delmotte, P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou. Cambridge University Press.

IPCC (2021b). “Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.” In *Climate change 2021: the Physical Science Basis*, edited by Valérie Masson-Delmotte, Panmao Zhai, Anna Pirani, Sarah L Connors, Clotilde Péan, Sophie Berger, Nada Caud, Y Chen, L Goldfarb, MI Gomis, *et al.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. 2391 pages.

Krusell, Per, , and Anthony Smith (2020). “Climate Change Around the World.” Working paper, Stockholm University.

Nordhaus, William D. and Andrew Moffat (2017). “A Survey of Global Impacts of Climate Change: Replications, Survey Methods, and a Statistical Analysis.” Working Paper 23646, National Bureau of Economic Research.

Papageorgiou, Chris, Marianne Saam, and Patrick Schulte (2017). “Substitution between clean and dirty energy inputs: A macroeconomic perspective.” *Review of Economics and Statistics*, 99(2), 281–290.

www.bportugal.pt