IARIW 2024

IARIW 2024

Thursday, August 22 – Friday, August 30

# The Net Zero Transition and Aggregate Productivity

Emilien Ravigné

(University of Oxford)

François Lafond

(University of Oxford)

Paper prepared for the 38th IARIW General Conference
August 26-30, 2024

Session 2B-2, Productive and Inclusive Net Zero II

Time: Wednesday, August 28 , 2024 [16:00- 17:30 GMT]

# The impact of the Net Zero transition on aggregate productivity*

PRELIMINARY AND INCOMPLETE
COMMENTS WELCOME

Emilien Ravigné[1,2] and François Lafond[1,2]

[1]Institute for New Economic Thinking at the Oxford Martin School, University of Oxford
[2]Smith School of Enterprise and the Environment, University of Oxford

July 31, 2024

**Abstract**

Reaching net zero emissions requires deep transformations of the production system, necessarily affecting productivity. However, very little is known about the aggregate productivity effects of technological and structural changes implied by net-zero scenarios. We build a parsimonious and transparent framework to assess these impacts. Our approach uses price and output growth pathways available in net zero scenarios, and derives industry-level and aggregate productivity growth. In this draft, we present only our method.
Keywords: Aggregate productivity; Net-zero transition

## 1 Introduction

Most macroeconomic models assessing the trajectory to net zero emissions assume exogenous and constant productivity growth, making economic growth only marginally dependent on the technologies and policies implemented. To give a simple example, typical models predict that under net zero scenarios, GDP in 2050 will be a couple of points higher or lower than in scenarios without the transition, but in both of these scenarios GDP in 2050 is roughly 3/4 higher than today, driven by exogenous assumptions on productivity growth. In DICE, productivity growth is exogenous, at more than 5%/yr until 2100 (Barrage and Nordhaus, 2023).

Assuming exogenous productivity growth in climate models is problematic because transition scenarios typically require deep transformations of production systems, including through the creation and deployment of new technologies, which should almost necessarily affect productivity. Unfortunately, while there are now fairly detailed scenarios and plans for the net zero transition, laying out how various technological options are deployed, these scenarios do not provide an assessment of the impact on aggregate productivity.

*We gratefully acknowledge funding from UK Research and Innovation through Economic and Social Research Council grant PRINZ (ES/W010356/1), from Baillie Gifford, and from the Institute for New Economic Thinking at the Oxford Martin School, as well as comments from audiences at Oxford, the ECB Working Group on Economic Modelling, CSH Vienna, IEW, and EAEREE. contact: emilien.ravigne@smithschool.ox.ac.uk; francois.lafond@inet.ox.ac.uk.

Our key contribution is to build a parsimonious, transparent and replicable framework to translate published net zero scenarios into productivity growth. Noting that the net zero transition will affect directly and profoundly only a limited subset of industries, we evaluate the impact on these industries, and aggregate these using weights justified by a fairly general theoretical framework. More specifically, our method follows two steps.

First, we infer industry-level productivity changes implied by net zero pathways. While scenarios generally include paths of price and production for key technologies, it is not straightforward to translate these into projections of industry-level TFP. For instance, the price of a technology may decline simply because its inputs are cheaper, rather than thanks to productivity improvement. To solve this issue, we introduce a simple assumption - we posit that the share of productivity growth that is translated in higher volume rather than lower prices is constant. While this relationship does not necessarily always hold, and is not theoretically motivated, it can be tested in the data. This is the key assumption allowing us to translate technology-level price and output scenarios into productivity pathways.

Second, we aggregate these industry-level productivity shocks into aggregate productivity. Productivity shocks accumulate along the supply chain, thus the position of an industry in the production network (upstream or downstream) determines an industry's influence on aggregate productivity. However, we show that under fairly general conditions (and in particular, outside of an equilibrium framework), the growth rate of aggregate productivity is a weighted average of industry-level productivity growth rates, with weights given by "Domar weights", which sum up to more than 1. This procedure, known as "Domar aggregation" is extremely useful for two reasons: first, it allows us to deal with each industry independently, rather than having to model explicitly all the supply chain interdependencies between industries. Second, it allows us to consider industries as they are defined in the net zero scenarios, rather than as they are defined in typical industry classification systems. Domar aggregation theoretically holds only under specific assumptions, which is why, here again, we test the validity of this procedure on past data.

Because these two steps can be tested on past data, we verify that they would have led to almost unbiased forecasts if they had been used in the past to predict productivity based on future prices and output. Of course, these predictions would not have been perfect - therefore, we use the distribution of errors to construct prediction intervals, providing us with empirically validated predictions of aggregate productivity conditional on net-zero scenarios.

We are in the process of applying this method to the UK, using mostly UK Climate Change Committee (CCC) and International Energy Agency (IEA) scenarios, and comparing different cost estimates when available. Preliminary results suggest offsetting effects, with the renewable energy transition boosting productivity, thanks to cheap solar and wind energy, but hard-to-decarbonize industries (e.g. cement) contributing mostly negatively to productivity growth. We are still revising these results, and developing estimated for other industries.

**What is net zero?** The net zero transition is a frame of reference through which mitigation action against climate change is structured. It is a scientific concept acknowledging that the main parameter for climate is the greenhouse gas concentration in the atmosphere resulting from the balance of human-made and natural emissions and natural sinks (Fankhauser et al., 2022). Meeting the 2°C target with a 50% chance means the remaining global carbon budget as of January 2023 was 1200 $GtCO_2$ and 250 $GtCO_2$ for the 1.5°C target (Lamboll et al., 2023). With the current emissions level at around 40 $GtCO_2$/year, this means reaching net zero emissions around 2050, which supposes planning a long-term strategy.

Mitigation pathways feature a large set of heterogeneous measures: energy efficiency measures, the transformation of the energy generation sector, the adoption of new technologies and processes, often industry-specific, and the development of carbon removal techniques. The impacts of this multi-faceted transition are therefore difficult to anticipate and can vary greatly between scenarios but are likely to impact macroeconomic indicators, including productivity.

**Literature.** There are various streams of literature analysing the empirical and theoretical relationship between environmental policies and macroeconomic outcomes. Most of the literature on the relationship between environmental policies and productivity pre-dates the recent price parity achieved by renewable energy, and proxies environmental policies through higher energy prices.

The well-known "Porter hypothesis" (Porter, 1996) claims that more stringent environmental policies boost productivity through enhanced innovation. Meta-analyses in various countries from firm-level to country-level do not show robust and consistent results (Benatti et al., 2024; Cohen and Tubb, 2018; Ambec et al., 2013; Kozluk and Zipperer, 2015). The effects at the firm-level are generally heterogeneous, with the most productive firms exhibiting productivity gains while others see their productivity fall (Albrizio et al., 2017).[1]. Greenstone et al. (2012) finds that air quality regulations, the 1970 Clean Air Act, caused a 2.6% decline in total factor productivity (TFP) in the US. They modeled these TFP losses by attributing a non-productive and exogenous share of output to regulatory compliance activities (monitoring, reporting, and abatement). Marin and Vona (2021) find that increased energy prices have on average a negative short term impact on firm-level productivity. André et al. (2023) finds a negative short-term impact of energy price shock on productivity but a positive impact four-year after the shock. More recently, Colmer et al. (2024) find that the EU Emissions Trading System (ETS) induced regulated firms to reduce by 14-16% their CO2 emissions compared to non-regulated firm. They find no evidence of economic contraction, offshoring or carbon leakage. On the contrary, they find weakly positive effects on productivity, value added, investment, and employment.

A large literature estimates the impact of environmental policies – in particular through higher carbon or energy prices – on innovation, but does not quantify the subsequent impact on productivity (Aghion et al., 2016; Calel and Dechezlepretre, 2016). Directed-technical change models (Acemoglu et al., 2012) provide theoretical foundations on the impact of green policies – especially prices and subsidies – on R&D in the clean and dirty sectors. Substitution between factors leads to productivity impacts. Hémous and Olsen (2021) offer a literature review of these models and the empirical evidence of such phenomenon, implicitly highlighting the difficulty of calibrating such models, as it is difficult to obtain robust estimates of the various elasticity of substitutions.

Integrated Assessment Models have produced a large number of transition scenarios and productivity trajectories. DICE (Barrage and Nordhaus, 2023) and REMIND (Leimbach et al., 2010) assume exogenous productivity growth. WITCH (Bosetti et al., 2006) models endogenous R&D spending and their spillovers, based on a U-shape innovation frontier. IMACLIM is built on a bottom-up technology-grounded model fed with cost and productivity forecasts (Hourcade

[1] For a more thorough literature review of the impact of environmental policies on innovation see Popp et al. (2010).

et al., 2010; Crassous et al., 2006). However, the large number of assumptions, the complexity of the models and their great heterogeneity make it difficult to compare them and draw a clear message about the future of productivity under net zero pathways.

More directly related to our work, the recent paper by Hallegatte et al. (2023) proposes a modelling framework to incorporate net zero scenarios into a CGE model. Their motivation is very much like ours: making use of detailed techno-economic scenarios usually available at the country level, and avoiding equilibrium models that essentially assume that the transition will have a negative impact, by modelling it as a policy-induced distortion in an an otherwise efficient world. Their approach is to change the parameters of a CGE model to make it match the techno-economic scenarios, in different ways in different sectors. This has the advantage that having calibrated a CGE model, they can run counterfactuals.

The report by Bijnens et al. (2024) discusses the effects of climate policies on productivity using various methods. Overall, their main message, based on an "environmental" DSGE calibrated on various NGFS scenarios[2], is that an orderly transition (an early and gradual increase of carbon prices) is better for productivity than a disorderly transition (which assumes an initially too small increase in carbon prices, such that a sharp rise is needed later on to ensure climate targets). Despite huge differences in carbon price patterns, with the disorderly transition featuring prices close to 1200 euros, however, the differences in productivity remain small, around half a percentage point difference of labor productivity level in 2050.

In contrast to these approaches, our reduced-form approach is focused on gauging quantitatively the performance of the predictions, by evaluating the errors we would have made if we had used our approach in the past (i.e. "backtesting" or "hindcatsing"). Given time series of future prices and quantities, how good are our forecasts of aggregate productivity? Measuring this carefully naturally leads to making predictions with credible uncertainty ranges. Also, our approach is simpler and thus more transparent, as it requires calibrating essentially one equation, rather than numerically solving a complex model.

In the next section (2) we lay out our framework, describing the accounting framework, the motivation for Domar aggregation, the rationale behind our reduced-form assumption relating prices and quantities to productivity, and the statistical validation of the overall approach using hindcasting methods. Section 3 briefly concludes.

## 2 Theoretical framework

In this paper, we focus on quantifying the implications of net-zero scenarios on aggregate total factor productivity. TFP is the main long-run driver of economic growth, which itself is the main driver of income per person.[3]

[2] The Network for Greening the Financial System is a network of central banks developing scenarios for the net-zero and its impact, for use by financial institutions.

[3] In growth accounting, labor productivity growth can be decomposed into a contribution of TFP and a contribution from the growth of capital per worker, with each term usually contributing substantially. However, labor productivity can also be decomposed into a contribution of TFP (TFP growth divided the labor share), and a contribution of a change in the capital-output ratio (Fernald et al., 2017; Acemoglu, 2024). The capital-output ratio is usually quite stable in the long term, making it clear that the focus for understanding long term growth should be on TFP. In other words, in this view capital per worker is endogenous. In a Solow model, the capital-output ratio is stable and capital per worker depends on TFP. Here we focus on TFP, and in ongoing work we will characterise the impact of the Net-Zero transition

In a nutshell, our goal is to find a function that maps $\tau$ years ahead prices $p_{i,\tau}$ and quantities $X_{i,\tau}$ of various industries, as given by scenarios, into aggregate productivity growth between $t$ and $t+\tau$, $\hat{A}_{t,\tau}$. We solve this problem in two steps. First we map prices and quantities into industry-level TFP, using a reduced form relation we test empirically, and then we aggregate industry-level TFP into an aggregate number, using Domar weights, that is,

$$\ln\frac{A_\tau}{A_t} \equiv \hat{A}_{t,\tau} = f\Big(\{X_{i,\tau}\},\{p_{i,\tau}\}\Big) = g\Big(\{\hat{A}_{i,t,\tau}(X_{i,\tau},p_{i,\tau})\}\Big) = \sum_i \frac{\text{Sales}_{i,t}}{\text{GDP}_t}\Big(\zeta\hat{X}_{i,t,\tau} + (\zeta-1)\hat{p}_{i,t,\tau}\Big), \quad (1)$$

where we implicitly assume we know prices and quantities of all industries today. Eq. 1 is the main result of this paper. In the rest of this section, we explain how we arrived at this formula, and how to implement it, including how we calibrate the parameter $\zeta$.

### 2.1 Set-up

We extend the basic set-up of McNerney et al. (2022), which provides a framework showing how to aggregate industry-level productivity shocks into aggregate growth, based only on basic accounting relations. Here we extend it to multiple factors, to ensure that aggregate productivity is defined as total factor productivity, rather than simply labor productivity. We provide detailed derivations in Appendix A, and report a high-level description here.

Consider a closed economy with $N$ industries, $F$ factors, and $D$ final demand categories[4], where each industry $i$ produces a single homogeneous good. We assume that the Input-Output system respects the condition that, for each industry $i$, total sales equal total expenses

$$\underbrace{\sum_{j=1}^{N} X_{ji}p_i}_{\text{Intermediate sales}} + \underbrace{\sum_{l=1}^{D} C_{li}p_i}_{\text{Final demand sales}} = \underbrace{\sum_{j=1}^{N} X_{ij}p_j}_{\text{Intermediate expenses}} + \underbrace{\sum_{f=1}^{F} L_{fi}w_f}_{\text{Factor expenses}}, \quad (2)$$

where $X_{ij}$ is the amount of goods produced by $j$ and sold to $i$, $Z_{ij}$ is the flow of money from $j$ to $i$, $p_i$ is the price of $i$, $C_{li}$ is the amount of goods from $i$ purchased by final demand category $l$, $L_{fi}$ is the quantity of factor $f$ used by industry $i$, and $w_f$ is the price of one unit of factor $f$. We denote total quantities produced by an industry by $X_i \equiv \sum_{j=1}^{N} X_{ji} + \sum_{l=1}^{D} C_{li}$, and for any variable $X$ we denote its log growth rate of as $\hat{X} \equiv d\ln X/dt = \frac{dX/dt}{X}$.

### 2.2 Domar aggregation

Let us define industry-level Total Factor Productivity (TFP) shocks as

$$\hat{A}_i \equiv \hat{X}_i - \sum_{j=1}^{N} a_{ji}\hat{X}_{ij} - \sum_{f=1}^{F} \tilde{l}_{fi}\hat{L}_{fi}, \quad (3)$$

on the capital-output ratio.

[4] In ongoing work we are extending the framework to an economy that is open, and with a government that raises and redistribute a carbon tax.

where the $a_{ji}$ and $\tilde{l}_{fi}$ are the shares of intermediate inputs $j$ and factor $f$ in $i$'s total expenses, and aggregate productivity as

$$\hat{A} \equiv \underbrace{\sum_{i=1}^{N} \theta_i \hat{C}_i}_{\text{growth of GDP}} - \underbrace{\sum_{f=1}^{F} \kappa_f \hat{L}_f}_{\text{growth of factor use}},$$

where $\theta_i$ is the share of industry $i$ in the total value of final demand and $\kappa_f$ are the factor shares. In Appendix A, we show that

$$\hat{A} = \sum_{i=1}^{N} \lambda_i \hat{A}_i, \tag{4}$$

where

$$\lambda_i \equiv \frac{p_i X_i}{pY} \tag{5}$$

is the *Domar weight* of the industry $i$, that is, the ratio of the industry value of gross output, $p_i X_i$, to the value of GDP, denoted $pY$. Since GDP is the sum of industry-level value added, and industry-level gross output is typically greater than industry-level value added, the weights sum up to more than 1.

Eq. 4 is known as "Domar aggregation", after Domar (1961), working with a Cobb-Douglas production function in simple systems, suggested to use these weights to aggregate industry-level TFPs. It is also known as "Hulten's theorem", after Hulten derived the result in a more general setting. However, while in Hulten (1978) it is obtained from assumptions about production functions, cost minimization and general equilibrium, here we follow McNerney et al. (2022) and show that Domar aggregation follows simply from the accounting definitions[5], Eqs. 2 and 3.

Eq. 4 is extremely powerful because it allows us to aggregate industry-level productivity shocks while ignoring the details of the production network.

One limitation of this result is that it is derived for small changes. There may also be other reasons why it does not hold in practice, including issues related to data quality or index numbers. In practice, we test how this aggregation formula works, and keep track of the errors it produces to construct prediction intervals (Section 2.5 below).

### 2.3 Estimating industry-level productivity growth from prices and volume scenarios

Since Domar aggregation tells us how to aggregate industry-level productivity shocks, all we have to do now is to estimate these industry-level productivity shocks. However, estimating productivity growth from Eq. 3 is a daunting task, as it requires estimating the change in unit requirement for all inputs. Our key contribution here is to introduce a simple, transparent, and testable assumption that allows us to circumvent this problem.

[5] Hulten's result in a general equilibrium framework (Hulten, 1978) ignores that changes to relative productivity affect relative sector sizes. To take into account these reallocations, one would need to know how industries' output react to changes in relative prices and propagate the shocks through the whole network, leading to an aggregation equation that, generically, includes the production network and the values of the elasticities of substitution (Baqaee and Farhi, 2019). McNerney et al. (2022) show that the output multipliers fluctuate much less than the industry-level TFP terms.

First, note that net-zero scenarios typically include pathways for the prices and/or output of the industry of interest, $\hat{p}_i$ and $\hat{X}_i$. So we are looking for a way to predict TFP in the future using only future values of $x_i$ and $p_i$.

Let us start by noting that TFP growth can be written equivalently in its primal or dual form,

$$\hat{A}_i \equiv \hat{X}_i - \sum_{j=1}^{N} a_{ji}\hat{X}_{ij} - \sum_{f=1}^{F} \tilde{l}_{fi}\hat{L}_{fi} = -\left(\hat{p}_i - \sum_{j=1}^{N} w_j\hat{p}_j - \sum_{f=1}^{F} \tilde{l}_{fi}\hat{w}_f\right). \tag{6}$$

In plain English, total factor productivity growth can be seen as either the growth of volume, netted out of the volume growth of inputs, or as the decline in price, netted out of the decline in the prices of inputs. Introducing the notation

$$\hat{\Sigma}_{X_i} \equiv \sum_{j=1}^{N} a_{ji}\hat{X}_{ij} - \sum_{f=1}^{F} \tilde{l}_{fi}\hat{L}_{fi}$$

for the growth of the index of inputs of $i$, and

$$\hat{\Sigma}_{p_i} \equiv \sum_{j=1}^{N} w_j\hat{p}_j - \sum_{f=1}^{F} \tilde{l}_{fi}\hat{w}_f$$

for the growth of the index of prices of the inputs of $i$, the primal-dual identity of TFP, Eq. 6, can be rewritten as

$$\hat{A}_i = \hat{X}_i - \hat{\Sigma}_{X_i} = \hat{\Sigma}_{p_i} - \hat{p}_i. \tag{7}$$

Rearranging, this means that the growth of (nominal) sales can be written as

$$\hat{s}_i = \hat{X}_i + \hat{p}_i = \hat{\Sigma}_{X_i} + \hat{\Sigma}_{p_i}. \tag{8}$$

In other words, the growth rate of nominal sales is equal to the growth rate of the cost of inputs, plus the growth rate of the volume of inputs. We make the key assumption that the relative contribution of these two terms is constant, that is, we assume that there exists

$$\zeta = \frac{\hat{\Sigma}_{p_i}}{\hat{\Sigma}_{p_i} + \hat{\Sigma}_{X_i}}. \tag{9}$$

If $\zeta$ was indexed by industry and time, Eq. 9 would just be a definition. However, we assume that $\zeta$ is constant over time, country and industry[6]. Substituting assumption (9) into Eq. 7, and rearranging using (8) leads to

$$\hat{A}_i = \zeta\hat{X}_i + (\zeta - 1)\hat{p}_i. \tag{10}$$

Eq. 10 is extremely useful – it allows us to determine the growth rate of TFP of industry $i$ by knowing only its price and output growth – we do not need the price or volume growth of its

[6] We could have used industry-specific and/or country-specific $\zeta$. We prefer to use a universal value, because while it is possible that industry-specific values of $\zeta$ would lead to more accurate predictions, in practice this also introduces more noise so we face a classic bias-variance trade-off. This is also simpler and thus more transparent.

inputs. Eq. 10 means that productivity growth translates into output growth and price decrease (if $0 < \zeta_i < 1$) in fixed proportions.

Eq. 10 is a purely statistical relation, with little motivation on theoretical grounds; It is a reduced form relation, and it is not immune to the Lucas critique. However, this is not an issue here - we do not need to assume that this relationship is causal, and we do not proceed to any comparative static; we only need a statistical relation between price/output growth and productivity growth, allowing us to infer the productivity growth assumptions that are implicit in net zero scenarios, which only provide future prices and quantities. Appendix B elaborates on this point further.

Appendix E provides a thorough evaluation of the validity of Eq. 10 in past data using hindcasting methods, for a universal value of $\zeta$. That is, for various databases, we put ourselves in the past and assume that we know future values of prices and output, and try to predict productivity, just as we will do in the next section with net zero scenarios. Crucially, we show that we can find $\zeta$ such that Eq. 10 leads to almost unbiased forecasts. Furthermore, we use the forecast errors to characterise uncertainty in the next section. In other words, we will make predictions of productivity conditional on prices and quantities, and we will assume that our errors for these predictions are similar to the errors we would have made in the past if we had used this method.

### 2.4 Estimating industry-level productivity growth from the productivity growth of the dirty incumbent and the clean alternative

Typically, net zero scenarios describe a process whereby, at the industry level, a green technology replaces a dirty incumbent. Thus, we need to take into account the productivity impact of both the rise of the clean technology and the decline of the dirty technology. One approach could be to estimate Eq. 10 on the clean and dirty subsectors separately. However, this raises an issue with using Domar aggregation (Eq. 4), as the Domar weights of the clean and dirty subindustries are, by definition of a "clean transition", changing substantially, so we would need to update them dynamically.

Instead, we propose to consider the clean and dirty alternatives together within each industry. That is, we construct an aggregate price and quantity index for the industry (say "electricity" or "steel"), and use the price and quantity index directly in Eq. 10 to determine industry-level productivity growth.

This raises an important question: what is the best index number for doing this? We have found that classic index numbers (Paasche, Laspeyre, Fisher, etc..) can give pathological results, in the sense that substituting a cheap dirty technology by an expensive clean technology can give price declines and thus positive productivity growth (see Appendix C). The intuition for this is simple: if both prices are non-increasing, a classic price index will be non-increasing; however, in the case of the net-zero transition, we are sometimes substituting a cheap product by an expensive one, so we would expect the industry-level price to increase.

Fortunately, here, we can make the assumption that each sector has a single, homogenous output, even though it may arise from different production techniques. For example, the electricity sector produces MWh of electricity, whether they come from fossil fuels or renewables. Similarly for the steel sector, a ton of clean steel from electric furnaces replaces a ton of dirty steel produced using coal.

Therefore, we use what is known as the "unit value" index (Balk, 2012): we express clean and dirty production in the same units, and define aggregate volume output as the simple sum. The price index is then the average of output prices weighted by their volumes.

### 2.5 Statistical validation of the framework and estimation of uncertainty

To sum up, our approach proceeds as follows. First, we select a few key industries that are typically important in net zero scenarios. For each industry, we look for scenarios for the price and quantities of the clean and dirty alternatives. We then construct industry-level price and quantity indices using unit value indices. Next, we use Eq. 10 to retrieve industry-level productivity scenarios from prices and quantities. Finally, we use Domar aggregation (Eq. 4) to compute the aggregate productivity impact.

Both our reduced-form approach to obtaining industry-level TFP and Domar aggregation are subject to errors. Here we test both assumptions on past data, and show how to use the errors to derive uncertainty intervals for our predictions. Note that our goal is to estimate the aggregate productivity impact associated with specific net-zero scenarios; we take the scenarios as given, and predict aggregate productivity conditional on these scenarios; we do not seek to characterise the uncertainty inherent to the price and quantity paths in the scenarios, but only the uncertainty associated with the use of our method to translate these scenarios into aggregate productivity numbers.

First, recall that Eq. 10 relates price and quantity growth to TFP growth. In Appendix E, using several databases, we use this equation to make predictions for the growth rate of productivity at horizon $\tau$, conditional of prices[7] and production growth, and assuming a universal value of $\zeta$, that is

$$\hat{A}^{\text{pred}}_{i,t,\tau} \equiv \zeta \hat{X}_{i,t,\tau} + (\zeta - 1)\hat{p}_{i,t,\tau}, \tag{11}$$

where for all variables we denote long growth rates between t and $t+\tau$ as $\hat{x}_{i,t,\tau} \equiv \log(x_{i,\tau}/x_{i,t})$, and the subscript $i$ denotes a given industry in a given country. We explore various values of $\zeta$, and choose one that minimizes the bias of the forecast.

The bottom left panel of Fig. 2.5 shows the distribution of the errors of these predictions, that is $\hat{A}^{\text{pred}}_{i,t,t+\tau} - \hat{A}_{i,t,t+\tau}$. We find that the predictions are close to unbiased, although with fairly large errors. We use the distribution of these errors to quantify the uncertainty, that is, we use the empirical quantiles of this distribution to construct prediction intervals around our point forecast given by Eq. 11. The three top left panels show how this procedure would have worked on past data for the manufacturing industry in three countries, showing that indeed, while predicting productivity from prices and output is far from perfect, it is possible to make useful predictions as long as uncertainty is accounted for in an objective and data-driven manner.

We proceed similarly with Domar aggregation (Appendix D): If in a given year in the past we had known future industry-level growth rates, and initial Domar weights, would we have

[7] In the scenarios, future prices are given in today's value of the currency. Therefore, to test the framework on past data, we need to express all prices in the same base year. In practice, we use the value added price deflator at the country-level to compute "real" price growth at the industry-level.

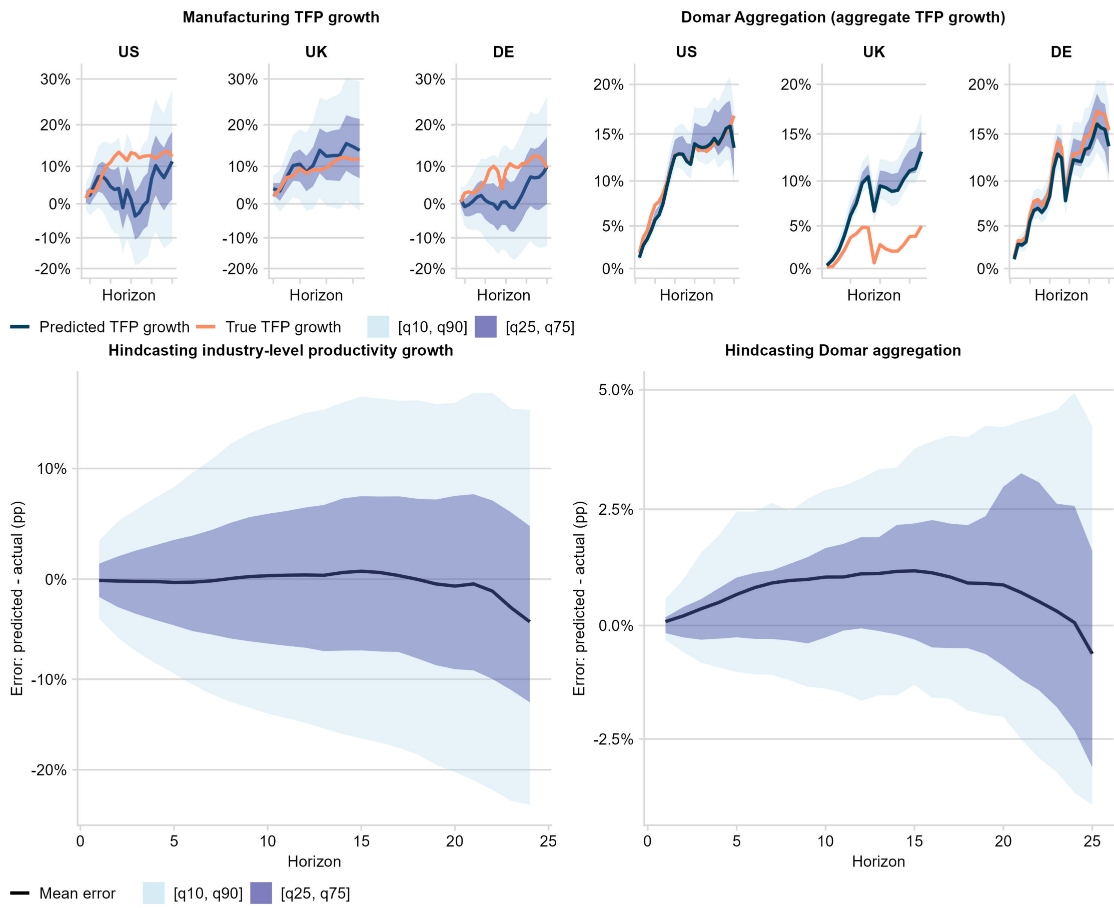


Figure 1: **Evaluation of forecast errors from the industry-level productivity growth predictions, and from Domar aggregation.** The left side of the figure presents the error in hindcasting the industry-level productivity growth with a value $\zeta = 0.128$. Top-left panels are hindcasting of TFP growth versus the true productivity growth in the US, the UK and Germany. Hindcasts are respectively in 1997-2020, 2002-2017, 1997-2020. The bottom-left panel shows the distribution of the error by horizon, pooling all countries and industries in KLEMS. The right side of the figure shows hindcasting results for Domar aggregation. Top-left panels are hindcasting of TFP growth versus the true productivity growth in the US, the UK and Germany. The UK performs poorly in the horizon, but it is an outlier compared to the rest of the database. The bottom-right panel shows the distribution of the errors by horizon, pooling all countries and industries in KLEMS. Both hindcasts are unbiased at horizon 25 years.

predicted aggregate productivity correctly? We predict

$$\hat{A}^{\text{pred}}_{c,t,\tau} = \sum_{j=1}^{N} \lambda_{j,c,t} \hat{A}_{j,c,t,\tau}, \tag{12}$$

for each country $c$, year $t$ and horizon $\tau$, using all industries $j$. The bottom right panel of Fig. 2.5 shows the distribution of the error, $\hat{A}^{\text{pred}}_{c,t,\tau} - \hat{A}_{c,t,\tau}$, showing a small positive bias for most horizons (at horizon 15 years, aggregate TFP growth is of the order of 5-10%, so a 1pp bias is not huge). The three panels on the top right again show the predictions for three different countries, using the quantiles of the distribution of errors to construct prediction intervals, and overlaying the underlying ground truth. Notably, Domar aggregation performs poorly for the UK using this dataset, but as shown in Appendix D this is an extreme case.

We construct final uncertainty ranges for the aggregate combining the distribution of errors for the industry level productivity growth estimation and the Domar aggregation. We use a Monte-Carlo method: for each run, we draw $N$ values from the distribution of industry-level errors, we aggregate them with their respective Domar weights, and then we draw one value from the distribution of aggregation errors and add it. We repeat the process 10,000 times to get a stable distribution of aggregate errors.

## 3 Discussion and conclusion

In this paper, we have developed a parsimonious and flexible framework to estimate the productivity growth impact of net zero scenarios.

To estimate the industry-level productivity shock from the price and output data available in scenarios, we have developed a reduced-form, data-driven approach. We have then shown how to aggregate these industry-level productivity shocks without relying on any specification of the production function or general equilibrium assumptions.

Our approach is transparent, tested in past data, and uses the errors made in testing to characterise the uncertainty of the predictions. We think this makes our results more robust than more sophisticated integrated assessment models, and in any case, provides a clear benchmark against which other methods should be compared.

In ongoing work, we are applying this framework to various industry-level scenarios, in key industries such as energy, transport, heat, steel and cement. Preliminary results indicate that (i) the aggregate effect is the result of somewhat counterbalancing industry-specific results, with energy likely featuring productivity gains and "hard-to-decarbonise" industries having productivity losses, and (ii) the aggregate impact is not linear in time, with costly upfront investment that pays off in terms of lower operating costs at longer horizons.

## References


Acemoglu, D. (2024). The simple macroeconomics of AI. Technical report, National Bureau of Economic Research.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The Environment and Directed Technical Change. *American Economic Review*, 102(1):131–166.

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Van Reenen, J. (2016). Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy*, 124(1):1–51.

Albrizio, S., Kozluk, T., and Zipperer, V. (2017). Environmental policies and productivity growth: Evidence across industries and firms. *Journal of Environmental Economics and Management*, 81:209–226.

Ambec, S., Cohen, M. A., Elgie, S., and Lanoie, P. (2013). The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Review of Environmental Economics and Policy*, 7(1):2–22.

André, C., Costa, H., Demmou, L., and Franco, G. (2023). Rising energy prices and productivity: Short-run pain, long-term gain? Technical Report 1755, OECD Publishing.

Balk, B. M. (2012). *Price and quantity index numbers: models for measuring aggregate change and difference*. Cambridge University Press.

Baqaee, D. R. and Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica*, 87(4):1155–1203.

Barrage, L. and Nordhaus, W. D. (2023). Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model.

Benatti, N., Groiss, M., Kelly, P., and Lopez-Garcia, P. (2024). Environmental regulation and productivity growth in the euro area: Testing the Porter hypothesis. *Journal of Environmental Economics and Management*, 126:102995.

Bijnens, G., Anyfantaki, S., Colciago, A., De Mulder, J., Falck, E., Labhard, V., Lopez-Garcia, P., Lourenço, N., Meriküll, J., Parker, M., et al. (2024). The impact of climate change and policies on productivity. *ECB Occasional Paper*, 2024(340).

Bosetti, V., Carraro, C., Galeotti, M., Massetti, E., and Tavoni, M. (2006). A World induced Technical Change Hybrid Model. *The Energy Journal*, SI2006(01).

Calel, R. and Dechezlepretre, A. (2016). Environmental policy and directed technological change: Evidence from the European carbon market. *Review of economics and statistics*, 98(1):173–191.

Cohen, M. A. and Tubb, A. (2018). The Impact of Environmental Regulation on Firm and Country Competitiveness: A Meta-analysis of the Porter Hypothesis. *Journal of the Association of Environmental and Resource Economists*, 5(2):371–399.

Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2024). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System. *The Review of Economic Studies*, page rdae055.

Crassous, R., Hourcade, J.-C., and Sassi, O. (2006). Endogenous Structural Change and Climate Targets Modeling Experiments with Imaclim-R. *The Energy Journal*, 27:259–276.

Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal*, 71(284):709–729.

Fankhauser, S., Smith, S. M., Allen, M., Axelsson, K., Hale, T., Hepburn, C., Kendall, J. M., Khosla, R., Lezaun, J., Mitchell-Larson, E., Obersteiner, M., Rajamani, L., Rickaby, R., Seddon, N., and Wetzer, T. (2022). The meaning of net zero and how to get it right. *Nature Climate Change*, 12(1):15–21.

Fernald, J. G., Hall, R. E., Stock, J. H., and Watson, M. W. (2017). The disappointing recovery of output after 2009. *Brookings Papers on Economic Activity*, pages 1–58.

Greenstone, M., List, J. A., and Syverson, C. (2012). The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing.

Hallegatte, S., McIsaac, F., Dudu, H., Jooste, C., Knudsen, C., and Beck, H. (2023). The macroeconomic implications of a transition to zero net emissions: A modeling framework. *World Bank Policy Research Working paper*, 10367.

Hémous, D. and Olsen, M. (2021). Directed Technical Change in Labor and Environmental Economics. *Annual Review of Economics*, 13(1):571–597.

Hourcade, J. C., Sassi, O., Crassous, R., Gitz, V., Waisman, H., and Guivarch, C. (2010). IMACLIM-R: A modelling framework to simulate sustainable development pathways. *International Journal of Global Environmental Issues*, 10(1/2):5–24.

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

Kozluk, T. and Zipperer, V. (2015). Environmental policies and productivity growth: A critical review of empirical findings. *OECD Journal: Economic Studies*, 2014(1):155–185.

Lafond, F., Greenwald, D., and Farmer, J. D. (2022). Can stimulating demand drive costs down? world war ii as a natural experiment. *The Journal of Economic History*, 82(3):727–764.

Lamboll, R. D., Nicholls, Z. R. J., Smith, C. J., Kikstra, J. S., Byers, E., and Rogelj, J. (2023). Assessing the size and uncertainty of remaining carbon budgets. *Nature Climate Change*, 13(12):1360–1367.

Leimbach, M., Bauer, N., Baumstark, L., and Edenhofer, O. (2010). Mitigation Costs in a Globalized World: Climate Policy Analysis with REMIND-R. *Environmental Modeling & Assessment*, 15(3):155–173.

Marin, G. and Vona, F. (2021). The impact of energy prices on socioeconomic and environmental performance: Evidence from French manufacturing establishments, 1997–2015. *European Economic Review*, 135:103739.
McNerney, J., Savoie, C., Caravelli, F., Carvalho, V. M., and Farmer, J. D. (2022). How production networks amplify economic growth. *Proceedings of the National Academy of Sciences*, 119(1):e2106031118.
Metcalfe, J. S., Foster, J., and Ramlogan, R. (2006). Adaptive economic growth. *Cambridge journal of Economics*, 30(1):7–32.
Popp, D., Newell, R. G., and Jaffe, A. B. (2010). Chapter 21 - Energy, the Environment, and Technological Change. In Hall, B. H. and Rosenberg, N., editors, *Handbook of the Economics of Innovation*, volume 2 of *Handbook of the Economics of Innovation, Volume 2*, pages 873–937. North-Holland.
Porter, M. (1996). America's green strategy. *Business and the environment: a reader*, 33:1072.

# A Domar aggregation

McNerney et al. (2022) derive Domar aggregation from a basic accounting framework. However, their derivations consider only a single factor, labor, so their equations typically express real quantities by deflating using the wage rate. To ensure that we interpret and measure TFP shocks properly, we extend the framework to include capital. We also provide a much more straightforward proof.

Starting from the industry-level accounting identity,

$$\underbrace{\sum_{j=1}^{N} X_{ji} p_i + \sum_{l=1}^{D} C_{li} p_i}_{\text{Total Sales, } p_i X_i} = \underbrace{\sum_{j=1}^{N} X_{ij} p_j + \sum_{f=1}^{F} L_{fi} w_f}_{\text{Total expenses}},$$

we can take time derivatives and rearrange to get

$$\hat{p}_i + \hat{X}_i = \sum_{j=1}^{N} a_{ji} (\hat{p}_j + \hat{X}_{ij}) + \sum_{f=1}^{F} \tilde{l}_{fi} (\hat{L}_{fi} + \hat{w}_f),$$

where the $a_{ji}$ and $\tilde{l}_{fi}$ are the shares of intermediate input $j$ and factor $f$ in $i$'s expenses. Rearranging, we have the definition of industry-level TFP as

$$\hat{A}_i \equiv \underbrace{\hat{X}_i - \sum_{j=1}^{N} a_{ji} \hat{X}_{ij} - \sum_{f=1}^{F} \tilde{l}_{fi} \hat{L}_{fi}}_{\text{TFP, primal approach}} = \underbrace{-\left( \hat{p}_i - \sum_{j=1}^{N} a_{ji} \hat{p}_j - \sum_{f=1}^{F} \tilde{l}_{fi} \hat{w}_f \right)}_{\text{TFP, dual approach}}. \tag{13}$$

GDP (in value) is the sum of final consumption

$$\text{GDP} = pY = \sum_{i=1}^{N} p_i C_i. \tag{14}$$

Taking time derivatives,

$$\hat{p} + \hat{Y} = \sum_{i=1}^{N} \frac{p_i C_i}{pY} \hat{C}_i + \sum_{i=1}^{N} \frac{p_i C_i}{pY} \hat{p}_i = \sum_{f=1}^{F} \frac{w_f L_f}{pY} \hat{L}_f + \sum_{f=1}^{F} \frac{w_f L_f}{pY} \hat{w}_f, \tag{15}$$

so that assuming a price index

$$\hat{p} = \sum_{i=1}^{N} \frac{p_i C_i}{pY} \hat{p}_i,$$

we have real GDP growth as

$$\hat{Y} = \sum_{i=1}^{N} \frac{p_i C_i}{pY} \hat{C}_i. \tag{16}$$

Now we *define* aggregate TFP growth as the growth of (volume) GDP minus the factor-share weighted growth of the volume of factors

$$\hat{A} \equiv \underbrace{\sum_{i=1}^{N} \frac{p_i C_i}{pY} \hat{C}_i}_{\text{growth of real GDP}} - \underbrace{\sum_{f=1}^{F} \frac{w_f L_f}{pY} \hat{L}_f.}_{\text{growth of volume of factors}} \tag{17}$$

We want to show that Domar aggregation is valid, that is

$$\hat{A} = \sum_{i=1}^{N} \lambda_i \hat{A}_i, \tag{18}$$

where the "Domar" weight $\lambda_i \equiv p_i X_i / pY$.

*Proof.* We want to show that the definition of aggregate TFP, Eq. 17, is equal to (18). We show this using the primal side of TFP. Substituting the LHS of (13) into (18) and replacing $\lambda_i$ by its definition, we need to prove

$$\hat{A} = \sum_{i=1}^{N} \frac{p_i X_i}{pY} \left( \hat{X}_i - \sum_{j=1}^{N} \frac{p_j X_{ij}}{p_i X_i} \hat{X}_{ij} - \sum_{f=1}^{F} \frac{w_f L_{fi}}{p_i X_i} \hat{L}_{fi} \right) = \sum_{j=1}^{N} \frac{p_j C_j}{pY} \hat{C}_j - \sum_{f=1}^{F} \frac{w_f L_f}{pY} \hat{L}_f. \tag{19}$$

Distributing the LHS into three sums and simplifying all the $p_i X_i / p_i X_i$,

$$\hat{A} = \underbrace{\sum_{i=1}^{N} \frac{p_i X_i}{pY} \hat{X}_i}_{(1)} - \underbrace{\sum_{i=1}^{N} \sum_{j=1}^{N} \frac{p_j X_{ij}}{pY} \hat{X}_{ij}}_{(2)} - \underbrace{\sum_{i=1}^{N} \sum_{f=1}^{F} \frac{w_f L_{fi}}{pY} \hat{L}_{fi}.}_{(3)} \tag{20}$$

Consider the third term, and permute the sums to get

$$(3) = \sum_{f=1}^{F} \frac{w_f}{pY} \sum_{i=1}^{N} L_{fi} \hat{L}_{fi} = \sum_{f=1}^{F} \frac{w_f L_f}{pY} \hat{L}_f, \tag{21}$$

where we have used $L_f \hat{L}_f = \sum_i L_{fi} \hat{L}_{fi}$. Similarly for the second term, permuting the two sums and using the fact that since $X_j = \sum_i X_{ij} + C_j$, $\sum_i X_{ij} \hat{X}_{ij} = X_j(\hat{X}_j - (C_j/X_j)\hat{C}_j)$, it becomes

$$(2) = \sum_{j=1}^{N} \frac{p_j}{pY} \sum_{i=1}^{N} X_{ij} \hat{X}_{ij} = \sum_{j=1}^{N} \frac{p_j X_j}{pY} \hat{X}_j - \sum_{j=1}^{N} \frac{p_j C_j}{pY} \hat{C}_j. \tag{22}$$

Substituting Eqs 21 and 22 back in Eq. 20, the first term of Eq. 22 cancels out the first term in Eq. 20, so that we are left with the RHS of Eq. 19. □

# B Discussion of the relation between prices, output and productivity

Here we elaborate on the empirical and theoretical content of our assumption relating industry-level prices and volume growth to productivity growth,

$$\hat{A}_i = \zeta \hat{X}_i + (\zeta - 1)\hat{p}_i,$$

Eq. 10 in the main text.

One way to think about this assumption behaviourally is that Eq. 10 simply describes the extent to which productivity growth is translated into either volume growth ($\zeta \approx 1$) or price declines ($\zeta \approx 0$), although we do not restrict $0 < \zeta < 1$. Eq. 10 also shows that, assuming $\zeta > 0$, productivity growth depends positively on output growth, as has been well documented for labor productivity under the name of the Fabricant-Kaldor-Verdoorn law (Metcalfe et al., 2006), or in the "learning curve" literature when considering the growth of *cumulative* output (see Lafond et al. (2022) for a discussion of the differences between the two).

To get further insights into the mechanistic implications of this assumption, we can transform Eq. 8 into

$$\hat{X}_i + \hat{p}_i = \frac{\hat{\Sigma}_p}{\zeta} = \frac{\hat{\Sigma}_X}{1 - \zeta}. \tag{23}$$

If the firm increases its nominal sales by 10%, Eq. (23) implies that the unit cost of inputs grew by $\hat{\Sigma}_p = 10\% \times \zeta$, *no matter whether the sales of the firm increased because of higher volumes or because of higher prices*. While we intuitively think that firms pass down higher input prices as higher output prices, or higher volume inputs purchases as higher volume output, these relationships are broken down here.

Another way to highlight an issue is this. Eq. (23) also implies

$$\hat{\Sigma}_X = \frac{1 - \zeta}{\zeta} \hat{\Sigma}_p, \tag{24}$$

which can be read as a price elasticity of input demand relation, making the assumption of a constant $\zeta$ feel somewhat more familiar. However, generally (but not necessarily) we expect $0 < \zeta < 1$, which would imply a *positive* price elasticity. This makes it very clear that $\zeta$ does not characterise a behavioural relation, but simply the fact that in the data we generally expect some inflation, and some growth in volume, so that an observed relation like (24) makes perfect sense.

Again, and to conclude, we regard this relation as a useful reduced-form regularity, and we do not think of $\zeta$ as a "deep" parameter. How an industry translates its productivity gains into higher production or lower prices is likely to depend on market characteristics, especially competition, the elasticity of demand, and the rate of innovation. However, we think these characteristics are unlikely to drastically change in the medium term despite the net zero transition, so that the reduced form relation should remain a useful regularity.

## C The index number problem and why we assume homogenous industry-level output

To see why we prefer to assume homogenous output than to use index numbers, consider a simple but pathological example (Table 2). Our economy features a transition from a dominant dirty sector to a prominent clean sector, and the clean sector is initially more expensive but is getting cheaper.

Table 1: Simple example of a two-good economy

| | Period 1 | Period 2 |
|---|---|---|
| *Quantities* | | |
| Dirty | 90 | 50 |
| Clean | 10 | 50 |
| Total | 100 | 100 |
| *Unit cost* | | |
| Dirty | 600 | 600 |
| Clean | 1500 | 1200 |
| *Sales* | | |
| Dirty | 54000 | 30000 |
| Clean | 15000 | 60000 |
| Total | 69000 | 90000 |

Because both prices are non-increasing, prices indices are non-increasing, even though the price of an "average unit" is increasing.

More precisely, with the numbers from Table 2, the Paasche price index would be $(50 \times 600 + 50 \times 1200)/(50 \times 600 + 50 \times 1500) = 0.86$, indicating a *decrease* of roughly 15% of the price, even though we have replaced 40% of the conventional production by the expensive clean good. A decrease in price would translate, in our framework, into an increase in productivity. We consider this a pathological outcome given that we have replaced a cheap technology by an expensive one.

We have evaluated in Table 2 various price and quantity indices and the basic issue remains, so we prefer to assume that the dirty incumbent and the alternative produce a homogenous product. This is known as the "unit value" index in the index number literature (Balk, 2012). This index assumes homogeneity between the units, so that the price index is the average price paid for one unit of good, and the quantity index is the total number of units.

Table 2: Price and quantity indices

| | Laspeyres | Paasche | Fischer | Törnqvist | Unit Value |
|---|---|---|---|---|---|
| Price index | 0.96 | 0.86 | 0.91 | 0.91 | 1.30 |
| Quantity index | 1.52 | 1.36 | 1.44 | 1.47 | 1 |
| Sales index | 1.30 | 1.30 | 1.30 | 1.33 | 1.30 |

***Note****: The indices are computed using the numerical values in Table 1, and the definitions of the price indices can be found in Balk (2012). Sales index is the product of the price and the quantity indices for Fischer, Törnqvist and the Unit Value Index. For Paasche and Laspeyres, the sales index is the product of the Paasche price index and the Laspeyres quantity index and conversely.*

## D Empirical testing of Domar aggregation

In progress.

## E Empirical testing of the reduced-form relationship between TFP, prices and quantities

In progress.

## F Computing aggregate error

In progress.