2

BANCO DE PORTUGAL ECONOMIC STUDIES

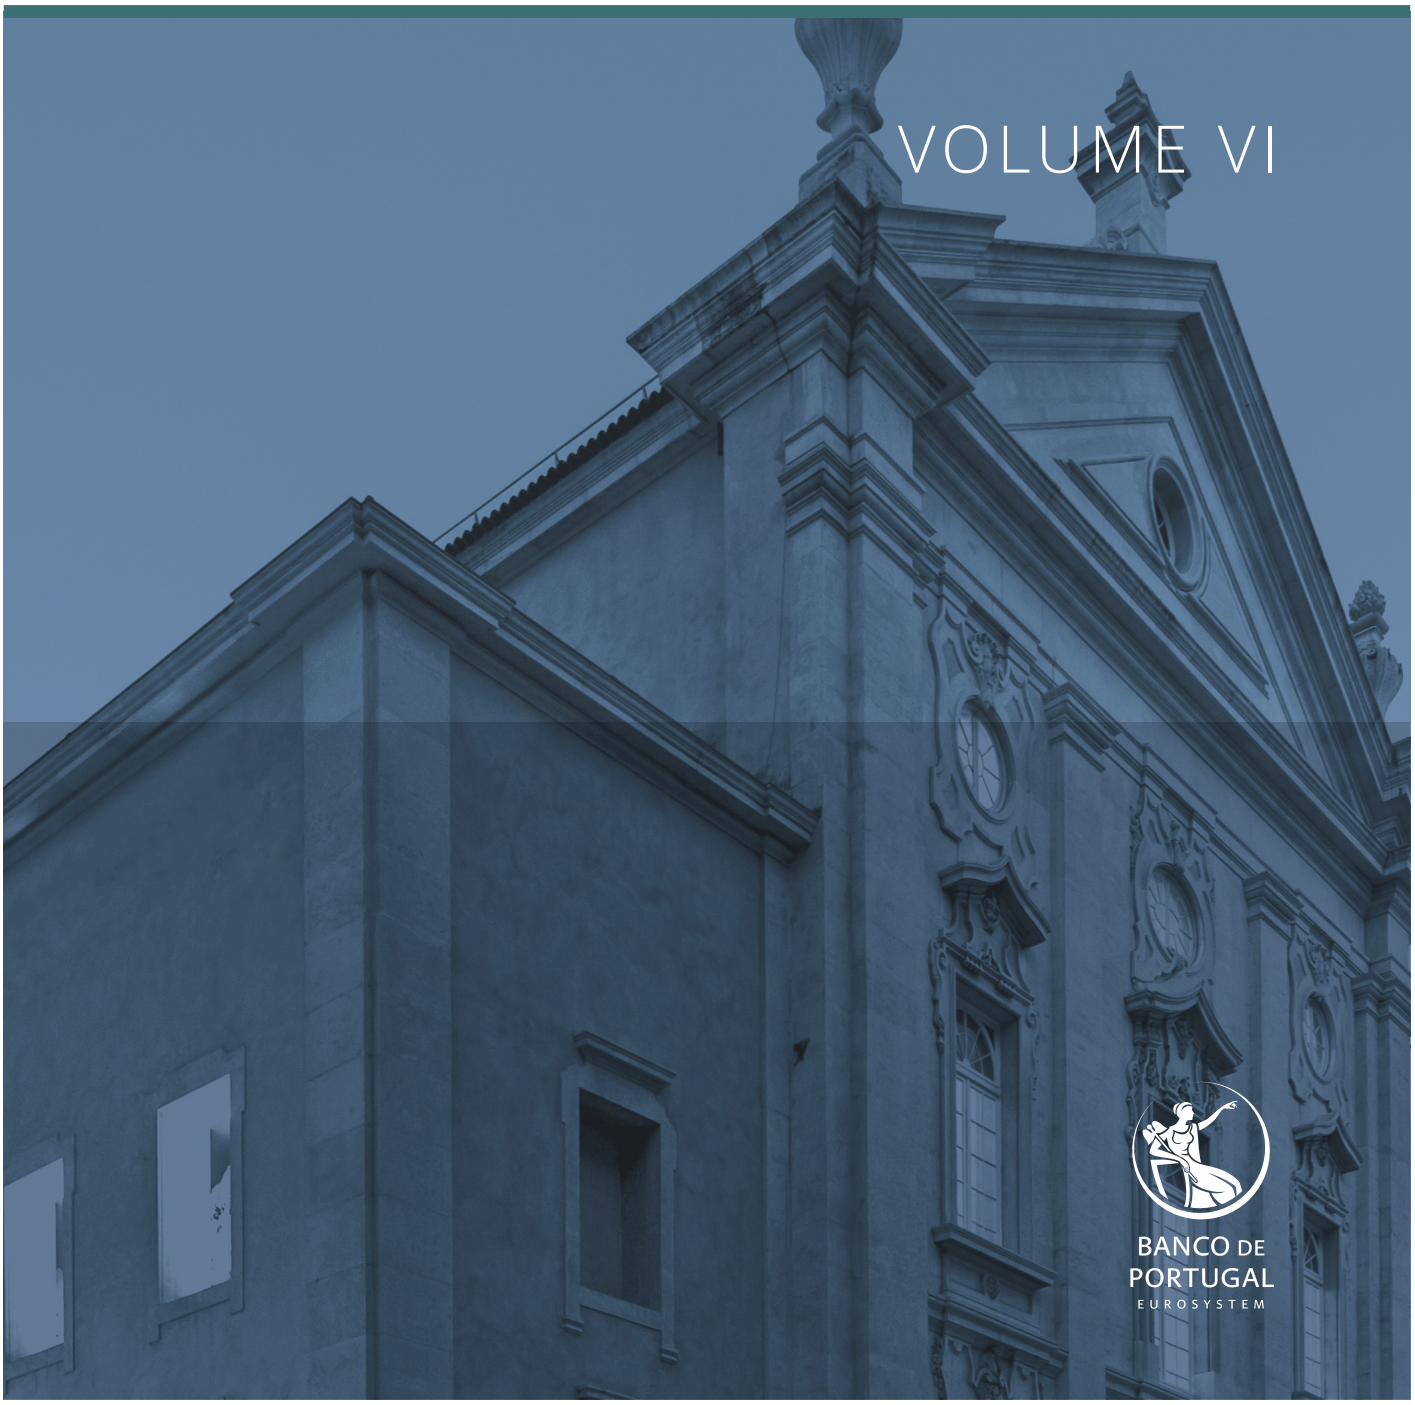

2

Banco de Portugal Economic Studies

Volume VI

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, 2020 • www.bportugal.pt

Banco de Portugal Economic Studies | Volume VI – no. 2 | Lisbon 2020 • Banco de Portugal Av. Almirante Reis, 71 | 1150-012 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

April 2020

---

1. Since the January issue of *Banco de Portugal Economic Studies*, the world economy has been suddenly affected by the COVID-19 pandemic. The spread of the pandemic and the ensuing social containment produced very extensive effects, which, despite the material economic policy response, are likely to lead to the largest post-war economic crisis. The truly unique nature of this global crisis results from the concurrence of supply and demand shocks which – in a context of utmost uncertainty and, therefore, extreme adjustments in economic agents' expectations – interact with one another, producing unprecedented magnifying effects[2] that generate recessional phenomena of an unusually large scale, in addition to effects in most commodity prices. As successive parts of the globe are affected by the pandemic, the crisis has a cross-cutting nature worldwide, which will, as projected by the International Monetary Fund, result in a reduction in per capita output in approximately 90 per cent of affected countries.[3] Finally, as in previous crises, this complex economic shock will tend to be amplified worldwide by adjustments through the financial channels, against a background where the crisis affects the various sectors of economic activity in an extraordinarily asymmetrical manner.

The unique nature of this economic crisis poses a new and particularly difficult challenge for the euro area. Despite the significant progress made since the Great Recession (which, in the euro area, ended up being a double-dip recession), the institutional framework of the economic area of which Portugal has been a member since its creation is unsatisfactory in at least three major aspects: the Monetary Economic Union is still incomplete, in particular because it does not have fiscal policy macroeconomic stabilisation mechanisms in place;[4] the Banking Union is also incomplete, as its three-pillar structure is not complete or in the process of being completed at all; and, finally, the

---

E-mail: pneves@bportugal.pt

1. The analyses, opinions and conclusions expressed in this work are entirely those of the editor and do not necessarily coincide with those of Banco de Portugal or the Eurosystem. This editorial was concluded in 15 April.

2. See, for instance, the Introduction to the March 2020 issue of the Economic Bulletin of Banco de Portugal and the e-book Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes, Baldwin, R. and di Mauro, B. W., A VoxEU. Org Book, CEPR Press, (2020).

3. On 9 April, Kristalina Georgieva said that "(. . .) we now project that over 170 countries will experience negative per capita income growth this year", in her speech called "Confronting the Crisis: Priorities for the Global Economy". See also IMF, World Economic Outlook, April 2020.

4. See, for instance, the presentation "A fiscal capacity for the euro area", Neves, P.D., at the "Economia Viva" conference, February 2020, available on Banco de Portugal's website.

Capital Markets Union has not made visible progress in the last decade, undermining possible gains in risk diversification among economic agents in the area and maintaining considerable pressure on the banking system, given its financial intermediation role, greater than in other economies. It remains the strong belief that, as in previous crises, the adverse circumstances will contribute to the deepening of the European project.

The decisions made by economic policy-makers, in particular at times such as these, require up-to-date information and economic indicators – both at the level of large macroeconomic aggregates and at the level of the various aspects of economic agent heterogeneity – and a thorough, in-depth knowledge of the functioning and adjustment mechanisms of the economy, likewise in aggregate terms as well as in terms of the distribution of behaviour and effects for individual economic agents. This is the challenge taken on by *Banco de Portugal Economic Studies*: to contribute to a better understanding of the Portuguese economy and to a better informed debate on economic policy.

2. This issue of *Banco de Portugal Economic Studies* includes four studies, focusing on varied subjects. The opening article describes the Exchange Rate Cooperation Agreement between Portugal and Cabo Verde, which, 20 years after its signature, has fulfilled its goals, thus becoming a reference case for international cooperation and the successful establishment of a fixed exchange rate peg. The other three articles are a good illustration of how the use of microeconomic data – covering, in three cases, time horizons of more than 20 years – makes it possible to strengthen the knowledge of how the Portuguese economy operates in the following domains: exports of goods, labour market *quantities* and labour market *prices*.

3. On 13 March 1998, in Praia, the Portuguese Minister of Finance and the Minister of Economic Coordination of Cabo Verde signed the Exchange Rate Cooperation Agreement between Portugal and Cabo Verde. From then on, the Cabo Verde escudo was linked to the Portuguese escudo – and, consequently to the euro as of 1 January 1999 – by a fixed exchange rate peg. In this way, an initiative by the authorities of the Republic of Cabo Verde has materialized, chiefly aiming at price stability and, more generally, macroeconomic and financial stability and the promotion of international openness. On that date, a body responsible for monitoring the Agreement was set up, in which representatives of the signatory ministries, ministries of foreign affairs and central banks of both countries participate. Banco de Portugal participated in the preparatory work for this agreement and, from the outset, in its technical follow-up.

A little over 20 years later, it is clear that the Agreement's objectives have been largely met. The exchange rate remained at the fixed exchange rate peg rate, inflation converged, on average, to values compatible with the definition of price stability, the economy experienced significant average real growth, while a pattern of gradual external openness prevailed. The positive balance of this Agreement was also acknowledged

by international institutions, such as the International Monetary Fund.[5]

The article by Correia, Cunha, Heitor, Maria and Saramago describes the Agreement, its main motivations and objectives, macroeconomic developments in Cabo Verde – in real and nominal terms, as well as regards public finances and external accounts – as well as some supply-side structural characteristics, such as demographics and the intensity of the use of labour. It also identifies a number of upcoming challenges to Cabo Verde's economy. This work gives therefore a very relevant and timely contribution to understanding the success of this Agreement, which, though not unique, is no less an interesting case study at international level.[6]

4. The use of microeconomic databases has been a predominant feature in Banco de Portugal's studies and economic analyses, since at least the early 1990s. The continuous availability of large-scale databases, covering – over time and across sectors – individuals, households and businesses, has been accompanied by remarkable progress in computing power constantly adjusting to developments in econometric theory. The analysis of heterogeneity in key economic variables has thus been crucial to understand the functioning and adjustment mechanisms of the Portuguese economy.

In addition to the databases developed and managed in-house,[7] the Bank has benefited from excellent cooperation with other national bodies, most notably, of course, Statistics Portugal. This issue of *Banco de Portugal Economic Studies* presents three studies exploring three microeconomic databases that have been particularly useful for the understanding of real and nominal variables of the Portuguese economy: exports of goods by firms located in Portugal, by Statistics Portugal; Statistics Portugal's Labour Force Survey; and "Quadros de Pessoal" (personnel records) of the Ministry of Labour, Solidarity and Social Security. In all these studies, it is worth highlighting the broad time-frame covered — 1993-2017, 1998-2019 and 1986-2013, respectively, i.e. periods always exceeding 20 years — which makes it possible to capture important structural aspects of the functioning of the economy and also its main cyclical features.

5. Export developments played a leading role during the recent expansionary phase of the Portuguese economy – now abruptly halted – decisively contributing to the

5. The IMF, for instance, states that: "The peg has served the economy well, given Cabo Verde's strong trade and financial ties to the euro area. (…) The fixed exchange rate regime has anchored macroeconomic stability in Cabo Verde". See IMF, Article IV Consultation – Staff Report, 25 September 2014.

6. For information on monetary agreements or foreign exchange-rate agreements involving euro area countries and the European Union and third countries, see "Monetary and exchange rate arrangements of the euro area with selected third countries and territories", ECB Monthly Bulletin, April 2006, and "Monetary and exchange-rate agreements between the European Community and Third Countries", Lamine, B., European Economy, Working Paper No 255, September 2006.

7. National and foreign researchers can access Banco de Portugal's microdata sets under predefined conditions via the Bank's microdata research laboratory (Laboratório de Investigação em Microdados – BPLIM). Also in this respect, the Bank contributes to the production of studies on the Portuguese economy.

current net lending position. The success of the Portuguese exporting sector in the period 2010-19 is illustrated in a clear-cut way by three economic indicators: (i) nominal growth of total exports by 72% (58% for goods, 103% for services – most notably the 152% increase in tourism exports); (ii) market share gains, which account for approximately half of the growth in exports by volume; and (iii) the share of total exports in GDP, which reached 44% in 2019, up by 16 percentage points from the average value observed between 1995 and 2007.[8] In addition, several studies and analyses have shown that this growth, supported by a wide range of products and services, was accompanied by gains in terms of technological content, innovation and differentiation and, overall, by a healthy demographic evolution of internationally active companies, where inflows and outflows have led to a selection by the market of the most competitive and efficient firms.[9]

The article by Cabral, Gouveia and Manteu presents ample evidence of an important feature of Portuguese exports of goods: their strong concentration on larger enterprises. Thus, in 2017, the 250 largest exporting enterprises accounted for almost 60% of total exports of goods (around 10 percentage points higher than in 1993), while the 1,000 largest exporting companies accounted for around 80% of total exports (5 percentage points more than in 1993). The high concentration of exports on the right tail of distribution – i.e. on a small number of (larger) enterprises – is highlighted in this study with recourse to three alternative statistical methods for the estimation of the Pareto exponent. Perhaps more interestingly, this study obtains the empirical result that idiosyncratic shocks affecting larger enterprises have a significant predictive power for total export fluctuations.

This study – by concluding that the developments in exports of goods of large enterprises and the specific shocks to which they are subject considerably help to explain the behaviour of aggregate exports – paves the way for the use of this information for very short term economic assessment and in short-term macroeconomic forecasts. As suggested by the authors, a more advanced use of this information could involve the use of a panel comprising the largest enterprises. It would be possible to anticipate, by some weeks, export data for these larger enterprises (the 100, 200, or 250 largest ones) and, thus, to identify possible idiosyncratic effects quantitatively relevant to the exports of goods; or even, if applicable, to identify at an earlier stage the turning points of the economy that result from the behaviour of external agents. Such a set of regular information could be of some use to provide an advanced quantitative indication of

8. The 1995-2007 reference period was established in such a way as to exclude the years 2008 and 2009, typically associated with the period known as the Great Recession.

9. See, for instance, the e-book Portuguese economic growth: A view on structural features, blockages and reforms, Banco de Portugal, Economics and Research Department, ed. João Amador, 2019, and the opening address at the 5th Central Balance-Sheet Database in Coimbra, Neves, P.D., 2016, Banco de Portugal.

developments in nominal exports of goods.[10]

6. The third study in this journal, by Martins and Seward, analyses the heterogeneity of the different states of the Portuguese labour market. Before highlighting the main contributions, it is worthwhile to provide some context on labour market developments in major economies in the period following the Great Recession.

At the time the world economy was affected by the exceptional dimension of the COVID-19 pandemic, the most advanced economies in general were in a position close to full employment of their inputs. Indeed, the unemployment rate had reached all-time lows[11] in several countries or economic areas at the end of 2019 or early 2020: 3.5% in February 2020 in the US, a 50-year low; 3.8% in December 2019 in the UK, a 45-year low; 5.5% in January 2020 in Canada, the lowest of the past 45-46 years; 2.2% in December 2019 in Japan, a 28-year low; 6.6% in January 2020 in the EU-27, the lowest since the start of this series in January 2000. The reduction in the unemployment rate to such low values naturally reflected the fact that, in 2019, the economies of the US, the UK, Canada and Germany had completed ten consecutive years of economic growth.

Such a marked reduction in the unemployment rate was not, however, accompanied by upward movements in wages[12] and prices that were compatible, in average historical terms, with such low levels of unemployment. The literature presents a broad set of reasons for this: price expectations anchored at levels compatible with price stability; low productivity that marked the recovery of advanced economies after the Great Recession; compositional effects on employment in terms of age, education and migration flows; effects of globalisation or technological change; weaker bargaining power of trade unions; or the so-called *pent-up wage deflation* according to which firms that face limitations to wage reductions during a recession (due to downward nominal wage rigidity) then tend to adjust wages slowly as the labour market recovers.[13] In

10. In this respect, Statistics Portugal and Banco de Portugal launched, in early April 2020, a survey of a wide range of small, medium-sized and large enterprises representing the various sectors of activity, with the purpose of identifying the effects of the COVID-19 pandemic on their activity (Fast and Exceptional Enterprise Survey – COVID-IREE). Statistics Portugal also publishes the twice-yearly Survey on the Perspectives of the Exports of Goods (IPEB) featuring the outlook of enterprises exporting goods about nominal export growth for year t (in November of year t-1 and May of year t). The relevant figures for this comparison are: 3.4% in May 1969 in the US; 3.8% in February 1975 in the UK; 5.4% in October 1974 in Canada; 2.1% in July 1992 in Japan.

11. The relevant figures for this comparison are: 3.4% in May 1969 in the US; 3.8% in February 1975 in the UK; 5.4% in October 1974 in Canada; 2.1% in July 1992 in Japan.

12. See for instance "Understanding low wage growth in the euro area and European Countries", Nickel, C. et al. (2019), ECB Occasional Paper Series, 232. This paper by economists of the European System of Central Banks documents and attempts to explain low wage growth over the 2013-2017 period and the reasons for registered growth to be systematically below the figures forecasted by the ECB, the OECD, the European Commission and the ECB Survey of Professional Forecasters.

13. FED Chair, Janet Yellen, used these words in her speech at Jackson Hole, in the Summer of 2014, speaking about the paper "Downward Nominal Wage Rigidities Bend the Phillips Curve", Daly, M.C. and Hobijn, B. (2014), Federal Reserve Bank of San Francisco WP 2013-08. Ms Yellen's statement on this point

this context, the assessment of the explanatory power of the Phillips curve – and in particular the possibility of it becoming flatter and displaying non-linearities in relation to the cyclical position of the economy – has become a central issue in macroeconomic modelling.[14]

Against this backdrop, "labour market slack" measures have been developed which calculate alternative metrics of labour underutilisation other than the unemployment rate, taking into account for example part-time workers willing to work longer hours and some sub-groups of inactive workers. The statistical bodies themselves – the Bureau of Labor Statistics[15] in the United States and Eurostat[16] in the European Union – regularly disclose alternative indicators of labour underutilisation. In addition, both the FED[17] and the ECB[18] have used these series for economic analysis.

By way of illustration, a relatively recent analysis in the ECB's Economic Bulletin[19] considered as spare capacity not being used and incorporated into the conventional concept of under-utilised labour the following: the *underemployed*, i.e. those who are working part-time but who would be willing to work longer hours; and *marginally*

---

was as follows: "First, the sluggish pace of nominal and real wage growth in recent years may reflect the phenomenon of "pent-up wage deflation" [see Daly and Hobijn] The evidence suggests that many firms faced significant constraints in lowering compensation during the recession and the earlier part of the recovery because of "downward nominal wage rigidity"–namely, an inability or unwillingness on the part of firms to cut nominal wages. To the extent that firms faced limits in reducing real and nominal wages when the labor market was exceptionally weak, they may find that now they do not need to raise wages to attract qualified workers. As a result, wages might rise relatively slowly as the labor market strengthens. If pent-up wage deflation is holding down wage growth, the current very moderate wage growth could be a misleading signal of the degree of remaining slack. Further, wages could begin to rise at a noticeably more rapid pace once pent-up wage deflation has been absorbed.' For empirical evidence on wage behaviour in Portugal see chapter by Martins, F. and Portugal, P. (2020) in the e-book Portuguese economic growth: A view on structural features, blockages and reforms, Banco de Portugal, Economics and Research Department.

14. See, for example, the following: "Phillips curves in the euro area", Moretti et al. (2019), ECB Working Paper Series 2295, who conclude that "… the Phillips curve is still a valid policy instrument once it is robustly estimated"; "A Phillips curve for the euro area", Ball, L. et al. (2020), ECB Working Paper Series, 2354, who state that a simple Philips curve captures most of the movements in inflation over the twenty years that the Euro has existed".

15. The U.S. Bureau of Labor Statistics releases six Alternative measures of labor underutilization, most notably the U3 (official unemployment rate) and the U6 (Total unemployed, plus all persons marginally attached to the labor force, plus total employed part time for economic reasons, as a percent of the civilian labor force plus all persons marginally attached to the labor force). The latest figures for these two series, as at March 2020, are 4.4% and 8.7% respectively.

16. See for example "Unemployment and potential additional labour force statistics", in Eurostat's website, published in October 2019.

17. See for example the "Hornstein-Kudlyak-Lange Non-Employment Index", on the website of the Federal Reserve Bank of Richmond and especially these economists' original work titled "Measuring Resource Utilization in the Labor Market", published in the Economic Quarterly, Volume 100, No 1, Federal Reserve Bank of Richmond.

18. See "Assessing labour market slack", Economic Bulletin, Issue 3, ECB, 2017.

19. See "Assessing labour market slack", Economic Bulletin, Issue 3, ECB, 2017.

attached workers,[20] which according to this study, are those who are not currently seeking work, despite being available (the "discouraged" workers), and those who are actively seeking work, but are not available to begin work immediately. Each of these two components corresponded at the end of 2016 to around 3% of the euro area's labour force.

Following the research that has been done in the past at Banco de Portugal[21] and benefiting from the enormous wealth of information contained in the individual data of the Labour Force Survey, published by Statistics Portugal on a quarterly basis, the article by Martins and Seward analyses the heterogeneity of the various states of the Portuguese labour market by estimating the transition rates between them through models based on Markov chains. The Labour Force Survey makes it possible to divide people who are not employed into three different groups: *unemployed*, inactive workers who want to work*(marginally attached)* and *inactive workers* who do not wish to or seek work. Based on the different rates of transition to employment, the study singles out five subgroups among the marginally linked individuals, highlighting the following results for each of them:

i) The transition rate of inactive workers not seeking employment to employment is close to the transition rate of the short-term unemployed (31.0% vs. 25.7%);

ii) The transition rates for inactive workers seeking employment and for inactive workers who claim other reasons (mainly pursuing an education, doing training and retirement) are close to the transition rate of the long-term unemployed (16.4% and 13.2%, respectively, vs.14.3%);

iii) The transition rates for discouraged workers and inactive workers claiming personal reasons are midway between the transition rates of the long-term unemployed and the inactive workers who do not wish – or seek – work (9.7% and 9.3% vs. 14.3% and 3.0%, respectively).

This study is therefore a very important contribution to understanding *quantities* in the labour market and, in particular, to a better understanding of the behaviour of some groups of inactive workers in this market.

7. "Quadros de Pessoal" (personnel records) – a database compiled annually by the Ministry of Labour, Solidarity and Social Security that covers all firms with at least one employee – are an unusually rich source for studying wages in Portugal. This database monitors workers throughout their career, as there is a unique identifier for each worker, firm, professional category and collective bargaining agreement. In

20. See also "Underemployment and potential additional labour force statistics", Eurostat, Eurostat Statistics Explained, October 2019.

21. See "Labour market heterogeneity: distinguishing between unemployment and inactivity", Centeno, M. e Fernandes, P.A., Boletim Económico do Banco de Portugal, March 2004 e "How to Measure Unemployment? Implications for the NAIRU", Centeno, M., Maria, J.R., e Novo, A., Boletim Económico do Banco de Portugal, Summer 2010.

addition to particularly precise information on the components of remuneration, the personnel records also have very complete information on employees' characteristics (like age, gender, education, qualification levels, occupational category, job tenure) and firms (sales volume, location, industry). The truly unique characteristics of this database allow those researching the labour market to carry out a thorough study and a very complete characterisation of wage behaviour in Portugal.

The article by Pedro Portugal, the fourth in this issue of *Banco de Portugal Economic Studies*, is a summary of a wide range of results obtained by the author and co-authors on wage behaviour in Portugal, which are brought together in a single text in a very convenient manner. Based on the information provided by the personnel records database – which, as the author describes, makes it possible to analyse (i) *who the worker is*, (ii) *where the work takes place*, (iii) *what the worker does*, and (iv) *with whom* – this article summarises the sources of changes in wages in the Portuguese labour market as a result of the heterogeneity of workers, firms and occupational categories. This study therefore makes a very important contribution to understanding labour market *prices*.

Pedro Portugal's article presents the main results of his research on wage behaviour in the Portuguese economy, covering several aspects: sources of wage change; pairing between worker and firm (in terms of wages and productivity); wage dispersion; wage sensitivity to the business cycle; gender wage gap; monopsony power in wage setting; the effect of a firm closing on wages; wage premium linked to unionisation; wage premium associated with education and costs of schooling in the workplace. More than providing a synthesis or highlighting the findings, it is justified to invite the reader to, without further delay, read and analyse this article "The sources of wage variability in Portugal: a binge reading survey".

Non-technical summary

April 2020

Exchange Rate Cooperation Agreement between Portugal and Cabo Verde: characterisation, developments and challenges after 20 years

Ana F. Correia, Vanda G. Cunha, Fernando Heitor, José R. Maria, Luís Saramago

The Exchange Rate Cooperation Agreement between Portugal and Cabo Verde (ERCA), in place for more than two decades, is a success case that deserves a deeper analysis. This article describes the agreement, its main goals — promotion of macroeconomic and financial stability and enhancement of economic openness to the rest of the world — and their reach. Some of the challenges that can be anticipated are also discussed.

On the eve of the Economic and Monetary Union's third stage, which brought about the euro, the authorities of Cabo Verde took the initiative of changing its exchange rate regime in order to make it more suitable to the structural characteristics of the economy. For that purpose, they established contacts with the Portuguese authorities, which led to the conclusion of the ERCA in March 1998.

Underpinning the agreement is the establishment of a fixed exchange rate between the Cabo Verde escudo (CVE) and the Portuguese currency (initially 0.55 CVE/PTE and after the implementation of the euro 110.265 CVE/EUR). The fixed exchange rate is supported by three main pillars, established in the ERCA:

i) The commitment of the Cabo Verdean authorities to implement economic policy guidelines supportive of the fixed exchange rate;

ii) The availability by the Portuguese Treasury of a credit facility to be used, if necessary, to reinforce the Cabo Verdean foreign exchange reserves; and

iii) The creation of a structure to monitor and provide technical support to the agreement, with representatives of both countries finance and foreign affairs ministries and central banks.

After 20 years, it's evident that the goals of the ERCA were largely achieved. The exchange rate remained unchanged and macroeconomic performance points to the broad fulfillment of the main goals. Monetary policy, aiming to ensure price stability, turned out to be credible and inflation converged to levels close to the ones observed

in the euro area (Figure 1). Trade openness also increased significantly (Figure 2), in the context of accelerated economic and *per capita* average income growth (Figure 3).

Among the future challenges it's worth mentioning the necessity to continue to promote sustainable economic growth, framed by balanced public and external accounts. To this aim, contribute not only budgetary consolidation measures but also structural reforms that enhance the investment in human and physical capital and the potential growth of the Cabo Verdean economy.

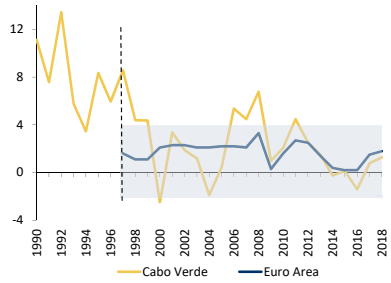

FIGURE 1: Inflation

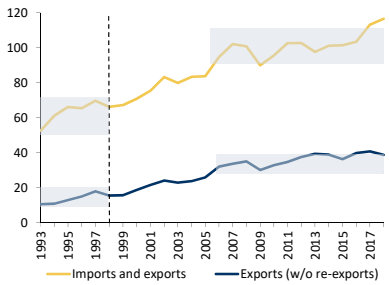

FIGURE 2: Degree of openness

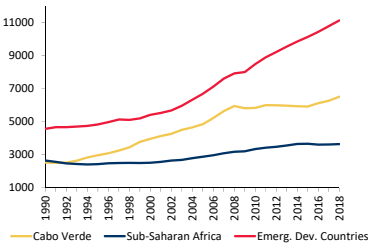

FIGURE 3: GDP *per capita*

Sources: BCV and IMF.
Notes: Inflation is measured by the CPI average rate of change, in percentage; the degree of openess is measured by the trade balance, in percentage of GDP; GDP *per capita* is in international dollars, using purchasing power parities of 2011.

Exchange Rate Cooperation Agreement between Portugal and Cabo Verde: characterisation, developments and challenges after 20 years

Ana F. Correia
Banco de Portugal

Vanda G. Cunha
Banco de Portugal

Fernando Heitor
Banco de Portugal

José R. Maria
Banco de Portugal

Luís Saramago
Banco de Portugal

April 2020

Abstract
The Exchange Rate Cooperation Agreement between Portugal and Cabo Verde (ERCA), in place for more than two decades, is a success case that deserves a deeper analysis. This article describes the agreement, its main goals – promotion of macroeconomic and financial stability and enhancement of economic openness to the rest of the world – and their reach. Some of the challenges that can be anticipated are also discussed. The macroeconomic performance points to the fulfillment of the envisaged goals. Monetary policy, aiming to ensure price stability, turned out to be credible and inflation converged to levels close to the ones observed in the euro area. Trade openness also increased significantly, in the context of accelerated economic and *per capita* average income growth. Among future challenges, it's worth mentioning the necessity to continue to promote sustainable economic growth, framed by balanced public and external accounts. To this aim, contribute not only budgetary consolidation measures but also structural reforms that enhance the investment in human and physical capital and the potential growth of the Cabo Verdean economy. (JEL: E31, E32, E42, H62, H63)

1. Introduction

Cabo Verde faced several challenges over the last decades. In 1998, the Ministers in charge of Finance matters in Portugal and Cabo Verde signed the Exchange Rate Cooperation Agreement (ERCA) that established two main goals: to promote the macroeconomic and financial stability of Cabo Verde and its greater openness

Acknowledgements: The authors would like to thank Nuno Alves, António Antunes, Rita Duarte, Malaquias Lopes, Sílvia Luz, Pedro Duarte Neves, José Carlos Pereira and Carlos Rocha for helpful suggestions and contributes. The analyses, opinions, and findings expressed in this article are those of the authors and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.
Note: This article was submitted for approval on February 21, 2020, before the World Health Organization declared the COVID-19 a pandemic on March 11.

E-mail: afcorreia@bportugal.pt; vmvcunha@bportugal.pt; fpheitor@bportugal.pt; jrmaria@bportugal.pt; lsaramago@bportugal.pt

to other countries, fostering economic and financial relations between Cabo Verde, Portugal and euro area countries.

The ERCA is built on the adoption of a fixed exchange rate peg between the currencies of both countries, which is supported by three main pillars, all inherent to the ERCA itself:

i) the commitment undertaken by Cabo Verde's authorities to adopt macroeconomic policies consistent with the regime;

ii) the provision by the Portuguese Treasury of a short-term credit facility at concessional terms, for the reinforcement of Cabo Verde's foreign exchange reserves; and

iii) the establishment of monitoring and technical support structures, composed of members of both countries' ministries of finance, ministries of foreign affairs and central banks.

After 20 years of ERCA functioning as envisaged, this article intends to present an analysis of the economic evolution of Cabo Verde in the light of the agreement's two main goals and to discuss some of the key challenges that lay ahead given the Cabo Verdean economy's structural characteristics. Macroeconomic performance points to the fulfillment of the established goals to a large extent. Monetary policy, mainly geared towards price stability, has proved credible, with inflation converging with euro area levels. As concerns the degree of openness, operations with non-residents increased, most notably with euro area countries. Regarding economic growth, though growth trajectories have started in the early 90's, average *per capita* income increased significantly in the decade that followed the agreement's signature, with Cabo Verde graduating from the UN's category of least developed countries (LDC) in December 2007.[1]

The future performance of Cabo Verde's economy depends on multiple factors. On the one hand, demographic change and higher average qualification levels should make a positive contribution towards a greater dynamism of economic activity. On the other hand, the economy faces a set of challenges to increase its growth potential. Structural reforms aimed at fostering investment (and consequently the capital stock) and increasing productivity levels are key to achieving higher sustainable welfare levels. Rising potential output would also help the budgetary consolidation process.

This article is organised as follows. In the first section some background considerations on the choice of exchange rate regimes are presented. The second section provides a description of the ERCA, its characteristics and goals. The third section analyses Cabo Verde's macroeconomic developments in the decade preceding ERCA's signature and in the two following decades, focusing on its goals. The fourth section discusses some of the Cabo Verdean economy's future challenges, namely demography and labour market, potential output and unemployment rate trend, for which estimates were computed. The article ends with some final considerations.

1. Cabo Verde was the second country ever to graduate, followed by other three since then (the LDC category currently includes 47 countries).

2. Brief background on the choice of exchange rate regimes

One of the key elements considered when designing the ERCA was the empirical and theoretical debate over the choice of the appropriate exchange rate regime, especially intense at the time, following the financial crises that hit emerging market and developing economies (EMDE). It should be remembered that since the collapse of the Bretton Woods system, in the early 70's, countries are free to choose the most appropriate exchange rate regime according to their goals: fixed, float or intermediate.

Typically, fixed exchange rate regimes involve the adoption of a foreign currency (e.g. dollarization) or the establishment of an explicit legal commitment to tie the country's currency official exchange rate to another country's currency (e.g. US dollar or euro) or to a basket of foreign currencies (currency board). In floating regimes, a country's exchange rate is set by the supply and demand relative to other currencies. Intervention may exist in the foreign exchange market to balance unjustified changes in the exchange rate of an ad-hoc nature (managed float) or mostly exceptional (free float). Between these two extreme cases, there are several intermediate regimes that assume different levels of commitment in relation to maintaining the national currency's peg to the foreign currency(ies) chosen. The objective is to narrow exchange rate volatility of floating regimes and, at the same time, keep some room for manoeuvre to conduct economic policy when responding to shocks.

The preferred exchange rate regime, especially for EMDE, has been changing over the last three decades. In the early 90s, pegging to a reference currency (usually the US dollar or the German mark) was a frequent choice, aimed at macroeconomic stabilisation. In the late 90s, after crises in the balance of payments in several EMDE, intermediate regimes lost support because they are hard to manage and more crisis prone, especially in the context of growing financial integration. The prevailing view on the choice of exchange rate regime pointed to either a fixed or a floating regime. This bipolar prescription is based on the hypothesis that in an open economy with price rigidity and free movement of capital, authorities aren't able to simultaneously have a fixed exchange rate, free capital movement and an independent monetary policy – the impossible trinity or trilemma (see Cruz-Rodriguez (2013)).

Nevertheless, after the collapse of Argentina's currency board in 2002 and faced with the fact that excessive exchange rate volatility can have negative impacts on these economies (fear of floating), the adoption of the bipolar prescription by EMDE has been questioned (Ghosh and Ostry 2009). Since then empirical and theoretical literature's advice over the choice of exchange rate regimes states that there is no "one size fits all" solution, i.e. there are no simple and unique prescriptions: the choice differs from country to country and, even for a specific country, it may vary over time. In this context, the choice of the most appropriate exchange rate regime depends on national authorities' goals and policy priorities, on restrictions imposed by the economic structure and/or functioning and on the country's institutional and administrative capacity.

3. Main Features of the Exchange Rate Agreement between Portugal and Cabo Verde

The Exchange Rate Cooperation Agreement between Portugal and Cabo Verde (ERCA) was signed on March 13, 1998, by the members of Government in charge of the Finance portfolio in the two countries.[2] It envisaged two main goals. On the one hand, to promote macroeconomic and financial stability in Cabo Verde, mostly by annulling the link between currency devaluation and inflation. On the other hand, to advance the economy's openness and the expansion of economic and financial relations between Cabo Verde and three layers of external partners: Portugal, the countries that would become members of the euro area (EA) and the countries with currencies pegged to the euro.

The agreement was part of a strategy focused on the liberalisation, modernisation and openness of the Cabo Verdean economy, aiming for the "dynamic insertion of Cabo Verde into the global economic system".[3] This strategy, which included several structural reforms, had among its highlights a privatisation program, the partial conversion of domestic public debt and the promotion of foreign competition. The broad-based privatisation program mostly encompassed financial sector firms (banks and insurance companies), together with water and electricity utilities, among others of lesser significance. The domestic public debt conversion was to be pushed forward through the issuance of new securities (*Títulos Consolidados de Mobilização Financeira*, TCMF). Their income would be linked to the yield of an off-shore Trust Fund built up with privatisation receipts and foreign partners' grants. This conversion would eventually encompass over half the domestic public debt stock, thus relieving the Cabo Verdean Treasury of its servicing. Measures taken to promote competition included such highpoints as the liberalisation of edible goods imports and the removal of administrative barriers to current account flows, as well as, in part, capital flows.

The initiative of changing the exchange rate regime came from the authorities of Cabo Verde. They approached the Portuguese authorities for that purpose, which led to technical meetings that involved the central banks and the ministries of finance, ultimately resulting in the ERCA. The latter is based upon a fixed exchange rate between the currencies of Cabo Verde and Portugal (at first 0.55 CVE/PTE and, after the euro was launched, 110.265 CVE/EUR). This decision also reflected the Cabo Verdean authorities'

2. This was the second experience of a peg between the currencies of Portugal and the Portuguese-speaking African countries following the latter's independences. The first had been the Monetary Arrangement between Portugal and Guinea-Bissau (1991/1996), which was not sustainably carried out but yielded important lessons (Saramago 2018). Later on, the Economic Cooperation Agreement between Portugal and São Tomé and Príncipe (ECA) was signed and has been functioning since 2009. It does not have the same explicit monetary and exchange rate dimension of the agreement with Cabo Verde, but it was inspired by the latter, and they are quite similar in their nature and functioning. Moreover, the ECA was set in parallel with a peg to the euro unilaterally established by the Santomean authorities (see, for instance, Saramago and Heitor (2015)).

3. See Governo de Cabo Verde (1997).

preference for a sort of exchange rate regime that was considered more suited to the structural features of Cabo Verde (Saramago 1998).

Among such features were, in particular, its condition as a small open economy (international trade in goods and services was 63.8% of GDP in 1993/1997), with a relatively narrow, rigid and little diversified production structure. It was, therefore, susceptible to the negative effects of currency devaluation (notably the pass-through effect on domestic prices) and unlikely to reap significant gains from a hypothetical positive impact of currency devaluation on external competitiveness. Another important feature was the high concentration of its relevant international transactions in Portugal and the future EA countries – either as partners in trade of goods and services or as the source of emigrant remittances, official development aid or foreign direct investment.[4]

Moreover, the financial system was relatively underdeveloped, in terms of sophistication, depth, diversity, or international connections. Such features, notwithstanding the progress achieved since the mid-1990s, might condition the potential benefits of opting for a flexible exchange rate regime, while at the same time curtailing the risks of going for a fixed exchange rate regime. Indeed, those features of the financial system might, on the one hand, hamper the exchange rate's effectiveness as a mechanism for adjustment. On the other hand, however, they might minimise the possibility of "speculative attacks" on the currency of Cabo Verde, along the lines seen in advanced or emerging market economies over the past few decades, in view of the limited significance, or absence, of financial markets and instruments.

The fixed currency peg is supported by three main pillars, set within the context of the agreement (as presented, for instance, in Saramago (2018)):

i) The first pillar was seen as the direct contribution of the Cabo Verdean party and it translated into a commitment to apply economic policy guidelines "compatible with the preservation of the exchange rate peg" – which meant preserving the macroeconomic foundations of the ERCA. This carried implications at several levels, as was then discussed. On the one hand, implications for monetary policy, which would need to be kept essentially aligned with the monetary policy stance adopted in Portugal and, after 1999, in the EA. At the same time, monetary policy would have to be even more focused on preserving international reserves, since their possible fall, if sizeable and sustained, would eventually imply the need for adjustments to the exchange rate. Concomitantly, among the functions of Banco de Cabo Verde, its main responsibility is to maintain price stability.[5]

On the other hand, as is also typical of fixed exchange rate regimes, the ERCA further carried implications for fiscal policy, which would have to refrain from deep

4. Note, for instance, that Portugal represented 40.2% of Cabo Verde's international trade in goods and was the country of origin for 18.3% of Cabo Verdean emigrants' remittances between 1993 and 1997. Considering all the six future euro area members of greater relevance as partners of Cabo Verde (EA-6: France, Germany, Italy, Netherlands, Portugal and Spain), those figures were 61.6% and 60.3%, respectively, in the same period.

5. In article 17, number 1, of Banco de Cabo Verde's Organic Law, as published in Boletim Oficial (July 15, 2002).

and prolonged expansionism. Otherwise, that might lead to external imbalances which, sooner or later, would put pressure on international reserves and thus on the ability to sustain the currency peg. There was also an understanding that the ERCA would be relevant for macroprudential policy (although the term was not used at that time yet). Indeed, safeguarding the financial system's stability was considered critical, as a first line condition to ensure the good functioning of the new exchange rate regime – against a background of significant changes, such as the main financial institution's privatisation and the entry of foreign capital into the financial system. And, to conclude, implications for structural policies as well, inspired by the notion that it would be important to provide an appropriate degree of flexibility to the functioning of goods and factors' markets, in order to compensate for the relative rigidity intrinsic to a regime such as the one underpinned by the ERCA.

ii) The agreement's second pillar was regarded as the direct contribution of the Portuguese party and it meant the Portuguese Treasury would provide a credit facility of a concessional and short term nature, sizeable in volume for the Cabo Verdean context (Table 1). [6] It was aimed at supplementing foreign exchange reserves in case of need – defined as a fall of the central bank's net foreign assets below a threshold equal to three months of imports of goods and services. The facility was used every year until 2004[7], after which international reserves remained above the threshold. It should be noted that the facility does not carry any obligation (of a monetary or financial nature) for intervention by Banco de Portugal or any other institution in the Eurosystem.

	1998	1999	2000
CF in months of imports (goods and services)	1.5	1.2	1.1
CF as a % of GDP	8.0	6.8	6.5
CF as a % of BCV net foreign assets	141.9	75.8	120.3

TABLE 1. Credit facility | Context

Source: BCV.

Note: The maximum amount available under the credit facility (CF) is considered in every year (i.e. €45 million).

iii) The third pillar was understood to be a joint contribution of both parties and it translated into the creation of a structure to monitor and provide technical support to the agreement. It was composed of the ERCA Commission (ERCAC) and the Macroeconomic Surveillance Unit (MSU). Joining the two parties on equal terms, this

6. This facility amounted to PTE 5.5 billion (5.5 million contos), rising to PTE 9 billion (9 million contos) against a collateral; after 1999: €27.5 million, up to €45 million. As collateral, a transfer of foreign exchange reserves was required, amounting to 35% of the requested value. Each withdrawal was set at PTE 1 billion (1 million contos); after 1999: €5 million. At a minimum, 80% of the outstanding amount would have to be reimbursed by the end of each year and the remaining (up to 20%) by the end of January in the following year.

7. There were three withdrawals in both 1999 and 2000, in each case for the maximum individual amount set under the ACC (€5 million). In the other years of that period (1998 and 2001-2004), there was a single annual withdrawal. In the years of 2001-2004, the Cabo Verdean party used the possibility of reimbursing 20% of the outstanding amount only by the end of January in each of the subsequent years.

structure brings together representatives of their ministries of finance, central banks and (in the case of ERCAC) ministries of foreign affairs. It enjoys privileged access to policymakers and has the technical capability to regularly perform an economic and financial assessment of the relevant context for the ERCA's functioning, identifying weaknesses and risks.

Launching the ERCA was directly connected to other components of the broader strategy undertaken by the Cabo Verdean authorities. It was, for instance, presumed that by lessening currency risk (and, more generally, by eventually raising predictability and confidence), the agreement might contribute to attract more foreign capital to the privatisation process. It was further expected that the nominal stability presumably fostered by the agreement might represent an added factor of credibility, thus inducing foreign partners to support the Trust Fund which would underpin domestic debt conversion.

Perhaps even more significant was the fact that the ERCA tended to be seen as a way to provide real content to the engagement between Cabo Verde and the European Union, which the Cabo Verdean authorities were openly trying to promote[8] (see, for instance, Correia (2002)). This engagement in the monetary and exchange rate field was further regarded as especially urgent in view of the fact that, while preparatory bilateral discussions were taking place (in 1997/1998), a deep structural change was about to happen – the start of the third phase of the Economic and Monetary Union (EMU), i.e. the launching of the euro on January 1, 1999.

It was largely assumed, by the time technical meetings between Portuguese and Cabo Verdean delegations were taking place, that Portugal would be among the early adopters of the single currency (as was later confirmed by the Council of the EU, on May 3, 1998). That prospect was then perceived to have at least two implications: i) the bilateral establishment of an agreement along the lines of the ERCA would become more difficult (if not impossible) within the framework of EMU; ii) in turn, if such a bilateral agreement were to be already operational at the beginning of the euro, it could be expected that European structures would adopt some kind of formal framework to accommodate the link between the Cabo Verdean escudo and the future common currency, as the successor of the Portuguese escudo.

This perception was grounded on the notion that the ERCA might receive a similar treatment to what was being prepared in the cases of the two CFA francs and the Comorian franc[9] — all of which maintained (for quite some time by then) fixed exchange rates against the French franc. Such pegs were supported by agreements among the African countries using those currencies as legal tender and France, another prospective

8. The disposition towards an engagement, which cut across the Cabo Verdean party-political spectrum, would lead to, among other relevant results, the "Special Partnership" agreed between Cabo Verde and the European Union in November 2007.

9. These are, respectively, the CFA franc issued by the Central Bank of Western African States and shared by the nine member states of the West African Economic and Monetary Union (WAEMU), the CFA franc issued by the Bank of Central African States and shared by five countries in Central Africa, and the national currency of the Comoros, an independent archipelago in the Indian Ocean.

founding member of EMU (see, for instance, ECB, 2006). And that was, indeed, what eventually happened. On December 21, 1998, the Council of the EU adopted a Decision which acknowledged the existence of the ERCA, involving a fixed exchange rate peg between the Cabo Verdean escudo and the euro. The agreement was kept on a bilateral basis but certain obligations of disclosure and consultation with the EU structures were introduced.

4. Macroeconomic developments in Cabo Verde (1990-2018)

How the ERCA functioned over the two decades since it was launched may be assessed by considering macroeconomic developments throughout those years, as compared to the previous period, with a special focus on the agreement's two main goals: macroeconomic and financial stabilisation, and economic openness. It is, naturally, important to bear in mind that even though the exchange rate regime represents a significant element to analyse such developments, it is nevertheless part of a much wider context of influences (Table 2). [10]

4.1. Product and income

Starting by how average income evolved in Cabo Verde, it becomes apparent that a fast pace of growth prevailed, starting before the ERCA, in 1993, and persisting until the global financial crisis of 2008. Over that period, real GDP *per capita* increased at an average annual rate of 5.5%, with some volatility but along a sustained trend (Figure 1a). Cabo Verde thus performed clearly better than two possible comparators: the whole of Emerging and Developing Countries (EDC: 3.4%) – within which emerging Asian countries tended to gain relevance – and, more pointedly, Sub-Saharan Africa (1.6%). Moreover, this performance of the Cabo Verdean economy took place against a background of significant population growth, at an average annual rate of around 1.7%.

This landscape changed significantly in the following decade, as real GDP *per capita* stagnated between 2008 and 2015, posting a recovery only in the three subsequent years[11]. The sizeable impact caused by the international financial crisis in Cabo Verde becomes therefore clear. Its growth path compares unfavourably with those of the two above mentioned comparators in the same period (2008/2015), and especially with the EDC, while being highly correlated with developments in the European economy (Figure 1b).

In the case of Cabo Verde, the crisis was felt mostly through the channels of tourism (gross receipts, at current prices, fell by 20.9% between 2007 and 2009, rising above the

10. This section benefitted from information contained in several reports of the two ERCA structures: the Macroeconomic Surveillance Unit (MSU) and the ERCA Commission (ERCAC), as well as in Banco de Portugal publications (namely "Evolução das Economias dos PALOP e de Timor-Leste") and IMF reports (in the context of Article IV consultations or reviews of programs established with Cabo Verde).

11. Note that population growth was slower in that decade too, but even so it reached an annual average rate of about 1.2%.

	Units	1993	1998	2003	2008	2013	2018
Output and prices							
Nominal GDP	CVE millions	34687	61679	92325	134698	153723	184661
GDP *per capita*	USD, PPP	1046	1377	1937	3639	3559	3579
Real GDP	Annual percentage change	7.5	8.1	4.4	6.7	0.8	5.1
GDP deflator	%	11.4	3.3	5.7	3.5	1.4	1.5
Inflation (annual average)	%	5.8	4.4	1.2	6.8	1.5	1.3
Inflation (end of period)	%	4.4	8.4	-2.3	6.6	0.1	1.0
External sector							
Current account	% of GDP	-5.7	-9.4	-10.2	-13.7	-5.9	-4.8
Capital account	% of GDP	4.4	3.0	2.8	1.5	0.3	0.8
Goods exports	% of GDP	0.9	1.7	6.5	6.4	9.9	13.8
Services exports	% of GDP	6.9	13.7	24.9	33.2	35.5	35.7
Tourism receipts	% of GDP	1.9	3.2	10.5	18.8	21.7	23.9
Goods imports	% of GDP	35.7	36.3	44.3	48.9	43.5	48.6
Services imports	% of GDP	4.4	14.4	23.1	19.5	19.1	19.5
Emigrants remittances	% of GDP	16.4	7.0	12.0	7.7	9.4	10.7
Foreign direct investment	% of GDP	0.8	1.4	4.1	11.6	3.0	4.3
Foreign exchange reserves	Months of imports G&S	4.3	1.1	2.0	4.0	4.8	5.6
Public finances							
Overall balance	% of GDP	-8.0	-3.3	-3.5	-2.1	-9.3	-2.7
Total revenue	% of GDP	32.5	27.4	23.9	28.7	24.5	28.2
Tax revenue	% of GDP	13.6	13.6	16.4	21.8	18.3	22.0
Grants	% of GDP	12.9	11.7	7.9	4.3	2.6	1.3
Total expenditure	% of GDP	40.4	30.8	27.4	30.8	33.8	30.9
Investment	% of GDP	23.8	13.5	8.8	9.2	10.8	4.3
Public debt	% of GDP	47.1	70.9	81.4	57.2	103.2	123.8
External	% of GDP	31.2	38.6	54.0	38.5	79.0	90.7
Domestic	% of GDP	15.9	32.3	27.4	18.6	24.2	33.1
Money and credit							
Net foreign assets	%	4.3	0.1	-1.7	-2.1	7.5	-2.2
Of the central bank	%	-4.9	-3.0	-0.5	2.4	4.5	0.7
Net domestic assets	%	20.6	2.6	10.3	10.0	3.5	3.9
Net claims on general government	%	7.8	-0.1	1.5	-1.6	2.0	0.5
Credit to the economy	%	-4.0	5.6	7.2	15.7	1.6	1.7
Broad money	%	24.9	2.7	8.6	7.9	11.0	1.7
Population and Labour market							
Resident population	Thousands	369.0	417.3	460.1	491.7	520.1	553.3
Unemployment rate	%	26.0	24.0	20.0	13.0	16.4	12.2
Employment	% total population	56.0	56.3	56.3	59.1	59.5	60.6
Labour force (15-64)	% total population	65.7	67.3	68.3	70.3	71.8	73.1

TABLE 2. Cabo Verde — Main economic indicators

Sources: BCV, Ministry of Finance of Cabo Verde, National Statistics Institute of Cabo Verde, IMF, World Bank.

Note: Values reported for Monetary Sector refer to changes as a percentage of initial M2.

previous peak, in nominal terms, only after 2011) and foreign direct investment (which had reached a peak amounting to 12.6% of GDP in 2007 and then fell almost every year, until the 3.0% of GDP recorded in 2013, recovering afterwards —- Figure 9b ahead). This more recent improvement of economic prospects in Cabo Verde benefitted also from the European economy's performance.

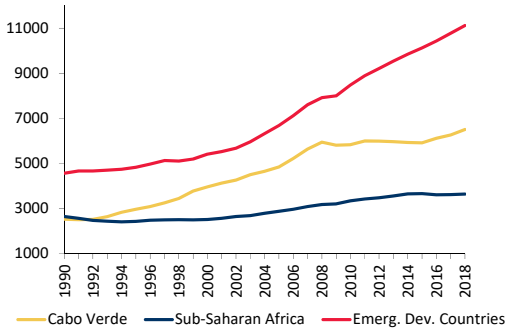

(A) GDP *per capita* | Constant prices

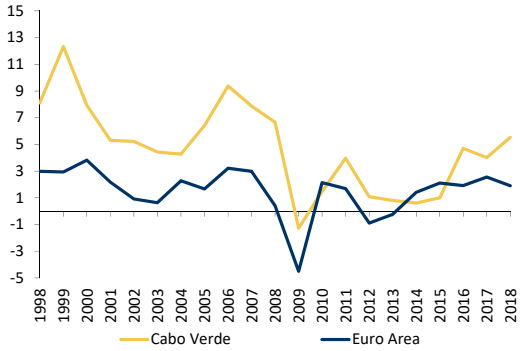

(B) Real GDP | Annual percentage change

FIGURE 1: Gross Domestic Product

Sources: BCV, Eurostat and IMF.

Note: GDP *per capita* is measured in International Dollars, 2011PPP.

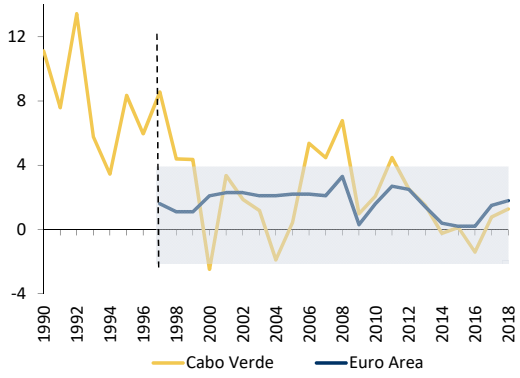

(A) Inflation | CPI, in percentage

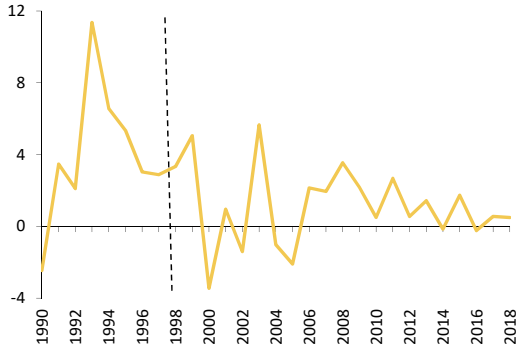

(B) Inflation | GDP deflator, in percentage

FIGURE 2: Inflation

Sources: BCV, Eurostat and IMF.

Note: Inflation is measured by the average annual growth rate.

4.2. Macroeconomic and financial stability goal

4.2.1. Inflation

Moving on to an assessment of Cabo Verde's performance in terms of macroeconomic and financial stability, one should start by looking at price developments, bearing in mind that the ERCA's currency peg would provide an anchor to promote the convergence of inflation levels in Cabo Verde with the lower ones prevailing in the EA. Looking at either the average annual change in the CPI (Figure 2a) or the annual change in the GDP deflator (Figure 2b), there is a quite clear contrast between the period of 1990/1997 and that of 1998/2018. In the first case, average inflation (using the CPI) stood at 8.0%, having remained almost always above 4%, while in the second case average inflation fell to just 1.9%, displaying a fast and sharp fall right in 1999/2000.

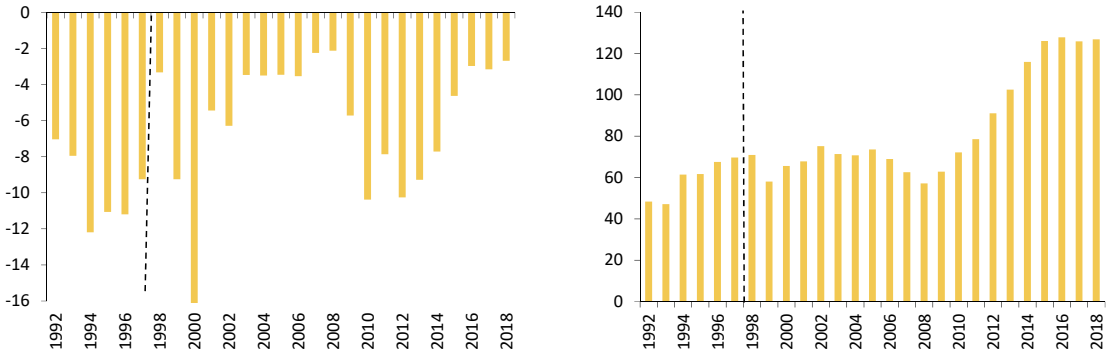

(A) General government balance

(B) Public debt

FIGURE 3: Public Finances

Sources: BCV, IMF and MF CV.

Note: Values as a percentage of GDP.

The pace was actually sustained from then on, in a consistent manner, with swings largely contained within a band of between -2% and 4%. There was just one more protracted rise above the upper threshold, in 2006/2008, of relatively low significance against the background of the whole period. There were also three spells of somewhat relevant negative inflation – in all cases linked to one-off shocks on the supply side (including good agricultural seasons, sectorial shifts favouring competition and changes in international prices of fuel, among others). The high level of convergence with the EA since the global crisis of 2008, in terms of CPI change, is also noteworthy.

4.2.2. Public finances

The analysis of public finance developments is also particularly relevant as fiscal policy matters in a fixed exchange rate regime and specifically in the context of Cabo Verde. In fact, the connection between fiscal and external imbalances may impact the foreign reserves level and the ERCA's sustainability itself. Since the signature of the agreement, two periods of high general government deficits with subsequent adjustments may be highlighted: in 1999/2000 and in 2009/2014 the general government deficit as a percentage of GDP has increased declining afterwards in both cases (Figure 3a). The first episode was related to the State assumption of SOE's liabilities in the context of their privatization[12] and to three exogenous factors to be mentioned: an exceptional drought period leading to the increase of social support expenditure; the USD appreciation *vis-à-vis* the Euro, implying an higher interest expenditure on the USD denominated debt; and the high international oil price translated into larger costs with subsidies to contain the domestic energy prices.

12. Note that the criteria for the treatment of SOE's debt assumptions in 2000, i.e. recorded as expenditure in national accounting, with an impact on the deficit, contrasts with the methodological treatment followed in the subsequent years.

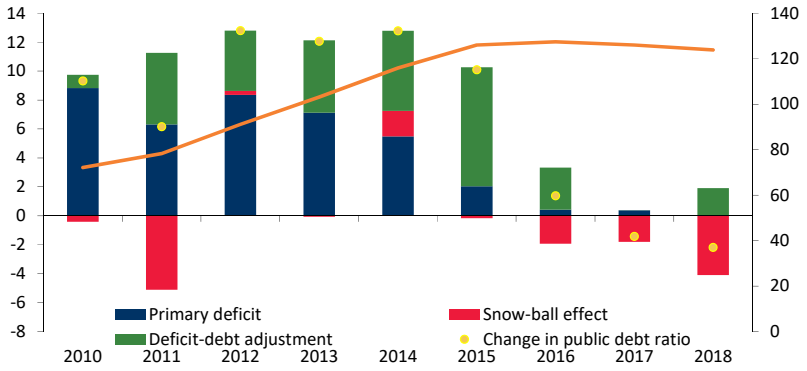

FIGURE 4: Public debt dynamics

Sources: BCV and authors' calculations.

Note: Values in percentage and percentage points of GDP.

However, this deficit increase did not mirror on the public debt ratio as the privatizations receipts were used to reduce the public debt (Figure 3b).[13]

In the second episode, that took place in 2009/2014, fiscal imbalances amplified, with an average general government deficit of 8.5% of GDP (with a peak of 10.5% in 2010) and the public debt ratio doubling from 57.2% of GDP in 2008 to 115.9% in 2014. These developments reflected the impact of the global financial crisis both directly (via drop in tourism revenues, for instance) and indirectly with the offsetting measures implemented by the Cabo Verde's authorities. In fact, the authorities adopted a package of stimulus measures aiming at mitigating the crisis impact on the economic activity and employment. By analysing the evolution of indicators like real GDP, GDP *per capita* or unemployment rate, this stimulus did not provide the aimed results. In contrast, the impact of this package was quite visible in the public accounts, notably in the public debt ratio, that showed an upwards trend until 2016 and remains in levels close to, but still above, 120% of GDP in the recent years (Figure 4).

A more detailed analysis shows that other factors beyond deficit have conditioned the public debt developments throughout this period. To this end, the change of public debt as a ratio of GDP (Δb_t) in each year may be decomposed as: i) the "snow-ball effect" i.e. the difference between the debt implicit interest rate (iir_t) eand the GDP nominal growth rate (g_t), multiplied by the previous year's debt as a ratio of GDP (b_{t-1}); ii) the symmetric of primary balance as a percentage of GDP (pb_t); e iii) and iii) deficit-debt adjustments as a ratio of GDP (dda_t). [14]

$$\Delta b_t = \frac{iit_t - g_t}{1 + g_t} b_{t-1} - pb_t + dda_t$$

Based on this decomposition, one may find three main factors for the evolution of the public debt ratio in the last decade: i) the primary deficits accumulation explains,

13. In this article, the public debt concept excludes the *"Títulos Consolidados de Mobilização Financeira"*.

14. The derivation of this decomposition may be found, for instance, in Amador *et al.* (2016).

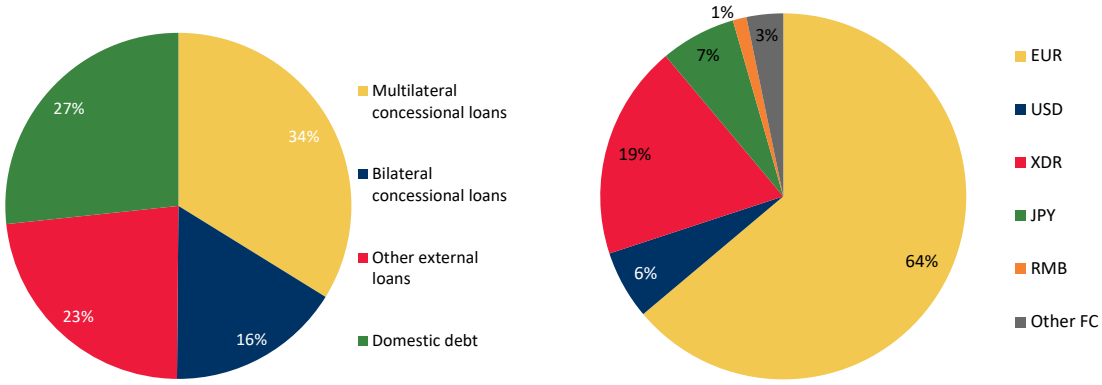

(A) By instrument | 2018 (B) By currency (external debt) | 2018

FIGURE 5: Public debt structure

Source: BCV.

Note: The values are in percentage of total.

at a large extent, the steep trend in the debt ratio up to 2014, only mitigated in 2011 by the macroeconomic performance; ii) the deficit-debt adjustments provided a significant contribution between 2011 and 2015, which has declined in the three following years. Two thirds of these adjustments refer to financing operations to State owned companies (SOE) through capital increases (not included in the general government balance) and on-lending operations which amounted to 23% of GDP in accumulated values from 2010 to 2018. Other factors impacting on public debt but not on the deficit as the exchange rate Euro/CVE depreciation in 2014 and 2015 had only a minor contribution to this item; iii) since 2016, the acceleration of economic growth along with budget consolidation measures carried out since 2014 have contributed favourably to the debt's dynamic given the positive differential between the GDP nominal growth rate and the implicit interest rate.

Analysing the Cabo Verde public debt structure at the end of 2018, some characteristics may be highlighted. By instrument, external concessional loans (multilateral and bilateral) represent half of the debt, and are typically long-term loans with low interest rates (around 1%) (Figure 5a). [15] Other external loans have contractual interest rates around 2%. Only the domestic debt issued in CVE has associated higher interest rates close to 5%. Breaking-down the external public debt by currency, about two thirds of the debt is denominated in Euro which, given the exchange rate agreement in place, narrows significantly the exchange rate risk exposure (Figure 5b). Additionally, almost 20% of the external debt in denominated in Special drawing rights (currency code XDR), and the Euro weighs close to 31% in this basket.

15. As Cabo Verde graduated from Least developed country status at the end of 2007, the concession of grants or concessional loans by external partners has been declining since then. The current concessionality profile of the external debt is not so favourable as in the pre-2008 period but this effect has been mitigated by the lower interest rates global environment.

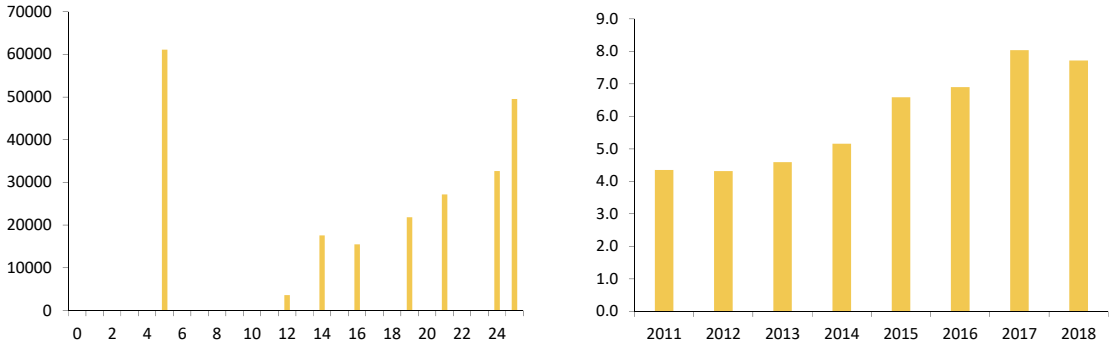

(A) By residual maturity, in years | 2018 (B) Sovereign guarantees to SOE

FIGURE 6: Public debt

Sources: BCV and authors' calculations.

Notes: public debt by residual maturity is in CVE millions. Sovereign guarantees to SOE as a percentage of GDP.

Decomposing by residual maturity, one may find a high average maturity and relatively distended from 2030 onwards. (Figure 6a). In the medium-term, a significant borrowing requirement is envisaged for 2023 corresponding to domestic loans denominated in CVE, which were already contracted with a high interest rate.

The evolution of contingent liabilities deserves to be observed as well. The sovereign guarantees to SOE's debt have exhibited an upwards trend, from about 4.5% of GDP in the early 2010 decade to around 8% of GDP in the last couple of years (Figure 6b).

Summing up, excluding the two periods of larger budget imbalances (1999/2000 and 2009/2014), associated as well to larger external imbalances, public accounts have revealed a relatively favourable performance since the signature of the ERCA. In fact, the years between 1992 and 1997 had witnessed considerable deficits of 9.8% of GDP on average. By contrast, the periods from 2000 to 2008 and since 2015 were marked by expressive budget consolidation processes, built both on revenues expansion associated to the economic recovery and, at some extent, on expenses control and improvements in the quality of expenditure.

Recent indicators on the institutional structure and governance confirm the rather favourable external perception of Cabo Verde. In particular, in the World Bank Governance Indicators (WGI) for 1996-2018 available for more than 200 countries, Cabo Verde position is around percentile 75 in the dimensions "voice and accountability", control of corruption and policy stability/absence of violence and terrorism and close to median in general government effectiveness and regulation quality. In the public sector Corruption Perception Index (CPI) compiled by Transparency International, Cabo Verde stays in the first quartile in 2019 (position 41 in 180 countries).[16]

16. World Bank (2019) and Transparency International (2020). The ranking for each indicator varies between 0 (lowest position) and 100 (highest position). Higher values correspond, in the case of WGI, to better governance, and in the case of CPI to a higher corruption perception level.

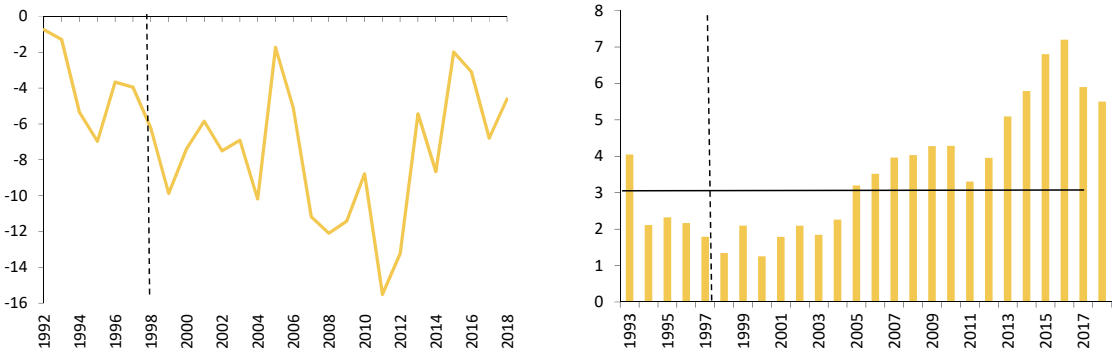

(A) Current and capital account | As a percentage of GDP

(B) Net external assets of BCV | In months of imports

FIGURE 7: External accounts

Source: BCV.

4.2.3. External accounts

As regards external accounts, a noteworthy feature is how the current and capital account deficit remained relatively high, as a percentage of GDP, for an extended period (Figure 7a). Among other reasons, this can notably be explained by the availability of external financing on favourable terms, in the shape of either foreign direct investment (Figure 9) or loans, namely to the public sector, on mostly concessional terms (which is to say, better than market terms).

Moreover, it is important to be aware of the structural change noticeable in the current and capital account, especially on the credit side (Figure 8a). The most remarkable fact is the growth of gross tourism receipts: amounting to no more than 2.5% of GDP in 1997, they leapt to 22.1% of GDP merely a decade afterwards.[17] They then took a toll from the global financial crisis, but would eventually reach a new peak in 2018 (24.0% of GDP).

In turn, emigrant remittances, still the major source of foreign currency in the 1990s (11.3% of GDP in 1996/1998), lost some ground over the following decade (8.7% in 2006/2008). They nevertheless recovered in more recent years, a trend that may be related to factors such as the economic recovery in the main countries of destination for Cabo Verdean emigrants and the falling yields on low risk savings in advanced economies since the global financial crisis. Lastly, official transfers (current and capital) – meaning official development assistance, in the shape of grants – lost relative

17. It is worth to notice that the tourism activity has a significant import content (given the narrow production basis of the Cabo Verdean economy vis-à-vis the tourism related demand). Therefore net increase of tourism recipts would be lower (although difficult to quantify as imports are not broken down by sectorial destination). An additional indicator of the growing importance of tourism in the Cabo Verdean economy is the GVA developments in tourism related sectors (accomodation and restaurants): it grew from 1.4% of GDP (average 1996/1998) to 6.3% of GDP (average 2016/2018).

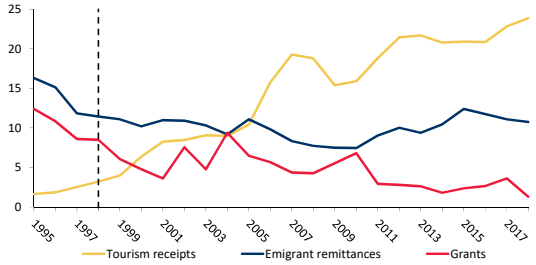

(A) Gross tourism receipts, emigrant remittances and grants

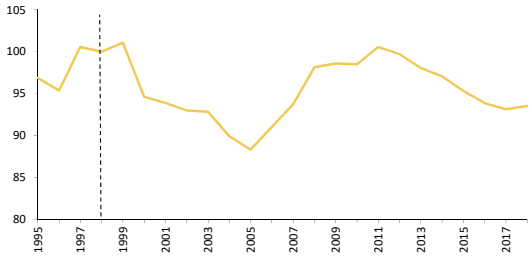

(B) Real effective exchange rate | Index 100=1998

FIGURE 8: Current account inflows and effective exchange rate

Sources: BCV, MF of Cabo Verde and INE of Cabo Verde.

Notes: Values in figure 8A are shown as a percentage of GDP. Tourism receipts correspond to gross values; the real exchange rate indicator was calculated by UAM using the bilateral rates for Cabo Verde's four main trading partners, weighed by their share in total trade of goods and deflated by their CPIs.

significance in a much sharper way, reflecting broad socioeconomic developments in Cabo Verde.

In order to better assess the soundness of Cabo Verde's external accounts, it is further necessary to look at how foreign exchange reserves evolved. Considering the BCV's net external assets, in months of imports of goods and services (Figure 7b), it becomes apparent that a downward trend prevailed in the first years after the agreement was introduced – reflecting an early adaptation process and the exogenous shocks mentioned in the context of public finances. In turn, from 2003 onwards this indicator strengthened almost continuously, except for the global crisis years, reaching a peak in 2016 (7.2 months of imports).

Still in the context of macroeconomic and financial stability, a final indicator of relevance is the real effective exchange rate of the Cabo Verdean escudo (Figure 8b) through which developments regarding price-competitiveness in the economy of Cabo Verde may be assessed. One of the main risks associated with fixed exchange rate regimes is, of course, the possibility that a relevant loss of price-competitiveness may occur as a result of gaps between domestic inflation and that of foreign partners. The risk increases with the size and persistence of such gaps, although one should bear in mind that overall competitiveness carries a non-price component too.

After some slight real appreciation, right in 1999, a period of six years followed, during which real depreciation tended to be continuous and sizeable. That gain in price-competitiveness would then diminish progressively, and it got totally wiped out by 2012, mostly as a result of higher inflation in Cabo Verde and the initial impact of the global financial crisis. It turn, the past seven years saw a return to the previous trend, of gradual CVE depreciation in real terms (an overall change of 7.0% between 1998 and 2018). It is therefore clear, according to the available indicator, that the euro peg regime did not imply a loss of price-competitiveness to the Cabo Verdean economy, taking into account the entire period since it became operational.

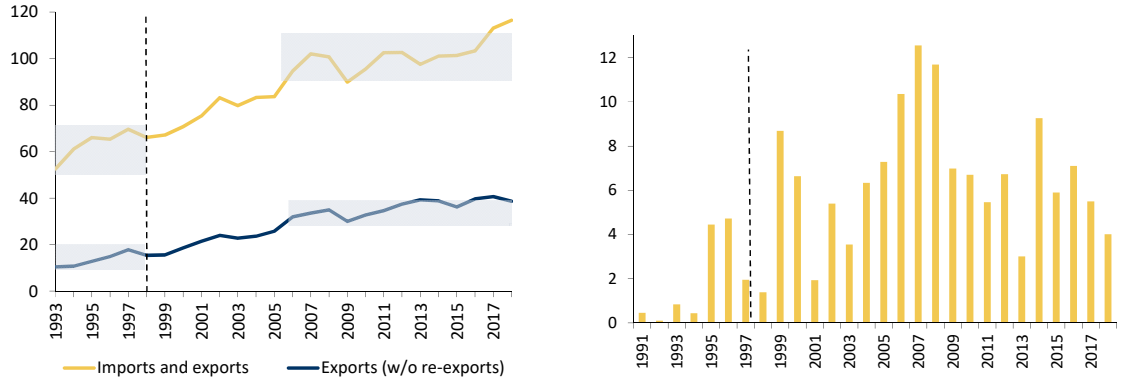

(A) Degree of openess | As a percentage of GDP

(B) Foreign direct investment (FDI) | As a percentage of GDP

FIGURE 9: Degree of openess and foreign direct investment

Source: BCV.

Notes: The degree of openess is measured by the level of external trade in goods and services. FDI comprises net inflows.

4.3. Economic openness goal

As regards the ERCA's second key goal – to promote greater openness in the Cabo Verdean economy and to foster its connections with the EA – it is important to start by identifying the rationale that underpins it: the lower uncertainty and higher confidence that were thought to result from abating currency risk towards the EA, which was by then already dominant in the context of Cabo Verde's economic relations with the rest of the world.

Considering the sum of exports and imports of goods and services as a percentage of GDP, a few results become apparent (Figure 9a). Following a period of relative stabilisation in 1995/1998, when it stood at around 65-70% of GDP, that indicator embarked on a sharp rising trend, clearing 100% of GDP about a decade afterwards. It then fell twice in the following years, which may be related to the initial impact of the global financial crisis and the beginning of the fiscal consolidation process. Nevertheless, it is clear that the degree of openness reached an all-time high in 2018 (116.4% of GDP) and has settled around 100% in the last decade (moving beyond that threshold in the past couple of years), in sharp contrast with the figures for 1993/1998.

This reasoning can be essentially replicated in case only exports of goods and services are considered – or rather, the trend appears to be sharper still. Actually, exports of goods and services more than doubled, as a percentage of GDP, since the ERCA was launched: starting at less than 18% of GDP before 1998, they would reach a peak of 40.6% of GDP in 2017. Indeed, the ratio of exports to imports evolved quite favourably: from an average of 28.1% of GDP in 1993/1998 to an average of 71.5% in 2013/2018 (peaking at 72.8% in 2018). It is therefore clear that the Cabo Verdean economy does not appear to have endured any loss of competitiveness, as regards its external trade, in the two decades since the ERCA came into force.

Still in the context of the external openness goal, it is important to take a look at foreign direct investment (FDI), comparing its performance before and after the ERCA was signed (Figure 9b). The underlying idea is that this sort of financing is especially sensitive to the level of confidence assigned by foreign investors to the destination economy's prospects (in this case, the Cabo Verdean one). And also the notion that an exchange rate framework such as the ERCA's might be regarded as an additional element of confidence, if it got to be credible.

The performance of FDI is actually quite suggestive: in the years between 1991 and 1998, it rose above 2% of GDP only twice. In turn, annual FDI fell below that threshold only once in the following couple of decades – and that happened in 2001, in the aftermath of the above mentioned critical phase. From 2002 until the global financial crisis, a consistent rise prevailed, with a peak in 2007 (12.6% of GDP). Even after 2009, FDI remained relatively robust (an average of 6.1% in the years between 2009 and 2018). It may thus be ascertained that FDI recorded a sizeable and permanent rise, as a percentage of GDP, in the period following the ERCA's introduction, broken only in times of international recession.

Moreover, it is relevant to consider how the weight of euro area countries in the context of Cabo Verde's external transactions evolved after the ERCA was signed. For the sake of clarity and data availability, exports (excluding re-exports) and imports of goods are considered, divided into three main groups: Portugal, the other five countries forming with Portugal the so-called EA-6 group (meaning the six euro area countries of greater significance in the context of Cabo Verde's external trade: Portugal, Spain, Italy, the Netherlands, France and Germany – in declining order of importance at the series' final year), and finally all the other foreign partners of Cabo Verde.

Even before the agreement was signed, Portugal was gaining ground since the beginning of the 1990s, from around 33-35% of total trade in 1991/1994 to 44.6% in 1997. That was decisive in pushing the whole of EA-6 countries beyond 60% of total trade in 1996/1998 (Figure 10). In the course of the next couple of decades, Portugal's weight never fell below 40%, having started by rising even higher (up to a peak of 56.2% in 2002) and tending to fall slightly afterwards – against the background of a growing presence of emerging countries such as China.

In turn, the five other EA-6 countries gained relevance until the global financial crisis (as the whole of EA-6 reached a peak of 85.3% in 2010, lower only than the all-time high of 2002: 86.5% of total trade – but, in the latter case, mostly due to Portugal). In a nutshell, the weight of EA-6 countries increased from an already noteworthy level of 59.7%, on average, in 1990/1997 to the rather more robust figure of 74%, on average, in the last few years (2012/2018), after remaining consistently above 80% in the period 2007/2011.

From the assessed indicators, it can be concluded that the Cabo Verdean economy was already on a path of rising openness since the beginning of the 1990s. Nevertheless, from the final years of that decade onwards, there was still a significant stepping-up in transactions with non-residents, especially those from euro area countries, as regards both current and capital flows (namely FDI, in the latter case).

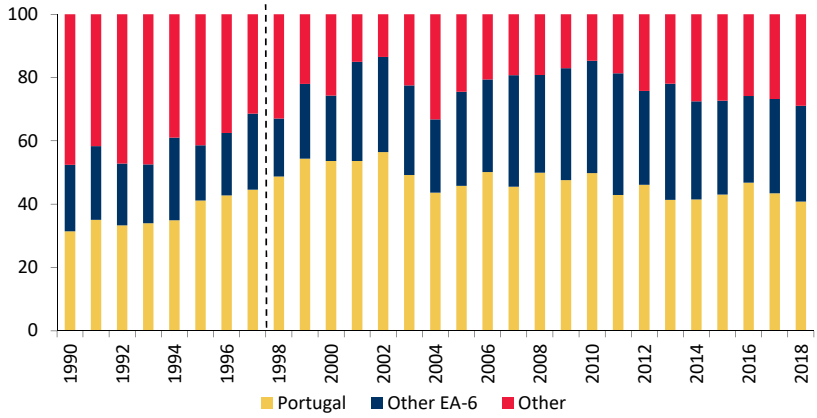

FIGURE 10: External trade, breakdown by countries

Source: BCV.

Notes: Data refers to exports (excluding re-exports) and imports of goods. EA-6 aggregate includes Germany, France, Spain, Italy, Netherlands and Portugal.

4.4. Key findings

Seeking to assess the impact that ERCA may have caused in the Cabo Verdean economy, the following points may be highlighted, as regards how the agreement functioned and how the economy broadly performed in the years 1998/2018 – see also Saramago (2018):

i) The exchange rate between the Portuguese and Cabo Verdean currencies remained unchanged since its initial fixing in the context of the agreement. As should be the case to safeguard confidence under this kind of exchange rate regime, preserving its effectiveness. It can thus be inferred that the existence and functioning of the ERCA never led to pressures so unbearable that an adjustment to the peg would become inevitable.

ii) The real effective exchange rate of the CVE appreciated only twice, in annual terms, against its level of 1998. Moreover, such phases were relatively mild and short-lived. For the whole period, the REER depreciated by 7.0%. Price-competitiveness, in as much as it can be ascertained through this indicator, was therefore preserved, avoiding one of the main pitfalls associated with this sort of exchange rate regime.

iii) The ERCA's credit facility was regularly used, in sizeable amounts, between 1998 and 2004, supplementing the Cabo Verdean foreign exchange reserves, as foreseen in the agreement. Since then, as the economy evolved, BCV net external assets remained above the threshold established under the ERCA as a condition of need.

iv) The technical structures created, on an even basis, to promote the agreement's functioning regularly carried out their responsibilities of macroeconomic monitoring. They thus analysed the ERCA's performance in the economic context of Cabo Verde, involving bilateral debates and the signaling of weaknesses and risks.

v) Inflation in the Cabo Verdean economy subsided consistently after the agreement was introduced. It thus converged with EA levels, in line with the first of the two major goals envisaged by the ERCA.

vi) External trade and foreign direct investment indicators show a significant increase in the openness of the Cabo Verdean economy to the outside world, which is in line with the second of those major goals.

5. Structural features and challenges of the Cabo Verde economy

5.1. Demography and labour market

Cabo Verde is a small lower middle-income economy[18] recording about 550 thousand inhabitants in 2018 but with a very young and increasing population. According to the latest Census, of 2010, when the resident population attained the half million, more than half of the residents were less than 25 years-old. The demographic pyramid notably shows the high density of the population in the youngest cohorts (Figure 11a). The young-age dependency ratio was about 50% and the old-age dependency ratio around 10%.[19] These figures contrast with the demographic structure of the European countries, including Portugal, where these ratios were 20 and 33%, respectively.

Since 2010 the resident population has been increasing by about 6,500 individuals each year, a lower growth than the one recorded in the decades prior to the crisis, but still a significant progression (corresponding to an average annual growth rate of about 1.2% - Figure 11b).

Note that the demographic developments mirror not only increases in the average life expectancy but mainly the still high fertility rates although diminishing. In fact, according to the results of the survey carried out by the INE and the Ministry of Health and Social Security (IDSR-III), the total fertility rate has declined from 2.9 in 2005 to 2.5 in 2018. Migration flows should have attenuated demographic growth via emigration, that has a historical tradition in Cabo Verde.[20] On the other hand, immigration has been increasing since the 1990s. Foreign population accounted for about 4% of total resident population in 2010 (3.2% in 2014).

The ongoing process of demographic transition initiated with the declining fertility rates is embedded in INE's demographic projections that point to a slightly lower growth of the population up to 2030. The projected annual growth rate is close to 1%, leading to 620 thousand individuals at the end of the horizon. The young-age dependency ratio is expected to drop to about 35% while the old-age dependency ratio

18. According to the World Bank classification. For purposes of the 2019's update, this classification applies to countries where the Gross Domestic Income *per capita* ranges between 1,026 and 3,995 USD.

19. The young-age dependency ratio corresponds to the weight of population younger than 15 over working-age population (ages 15 to 64), while the old-age dependency ratio corresponds to population older than 64 over working-age population.

20. INE CV has estimated about 16,500 individuals that had emigrated between 2009 and 2014, in particular young people that left to study abroad, and around 53% of the total moved to Portugal, followed by USA and France (Instituto Nacional de Estatística de Cabo Verde (2014, 2015)). To be referred that in 2018 about 34,660 individuals of Cabo Verdean nationality were resident in Portugal (Serviço de Estrangeiros e Fronteiras (2019)), which is equivalent to 6% of the total resident population in Cabo Verde in the same year.

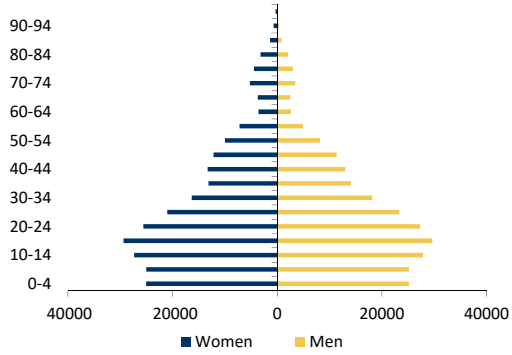

(A) Distribution by age and gender | 2010

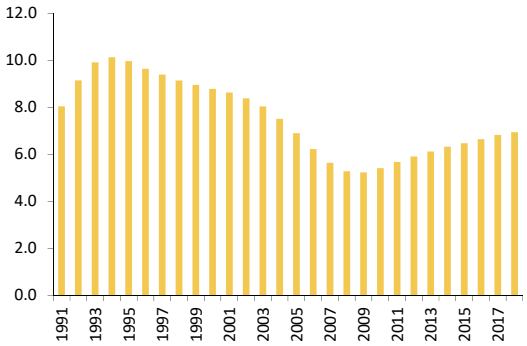

(B) Population | Change in thousands of indivudals

FIGURE 11: Population and distribution by age and gender
Source: INE CV (Censos 2010).
Note: Data for population is in number of resident individuals.

is projected to increase moderately to 12.5%. Overall, the total dependency rate should decline by then.

The increase of the working-age population and the higher labour market participation rates, notably female one, should continue to foster labour supply. In 2010, women's participation rate was about 56%, 15 percentage points below men's participation rate. However, until 2017, female participation rate grew above the male one (18.5% and 15.9% respectively). The unemployment rate has stood relatively high in the last decade (around 14%), but without relevant gender gaps. At this stage it should be noted that the informal sector is important in the economy of Cabo Verde, although this economy presents one of the lowest informality levels in the context of Sub-Saharan Africa. Declared employment is heavily concentrated in the services sector and, at lesser extent, in agriculture and fishing. This implies increased uncertainty in the statistical compilation of the labour market main variables. In any case, higher educational levels should allow for a more qualified labour force. In 2010 the increase of educational attainment levels was already quite evident in particular in the case of women (Figure 12b). This entails a significant increase in future average qualifications of the Cabo Verdean labour force.

5.2. Potential output and trend unemployment rate

The economic performance of Cabo Verde, like other economies, can be evaluated using unobserved variables, namely potential output and the unemployment rate trend component. The analysis of potential output level and growth rate adds information on the cyclical position of the economy and its sustainable growth.

Potential output is an unobserved variable that depends, among other factors, on the model that is used to estimate it. Duarte *et al.* (2019) suggest that the sign of the output gap is also conditional on the sample period and on the specifications of the unobserved

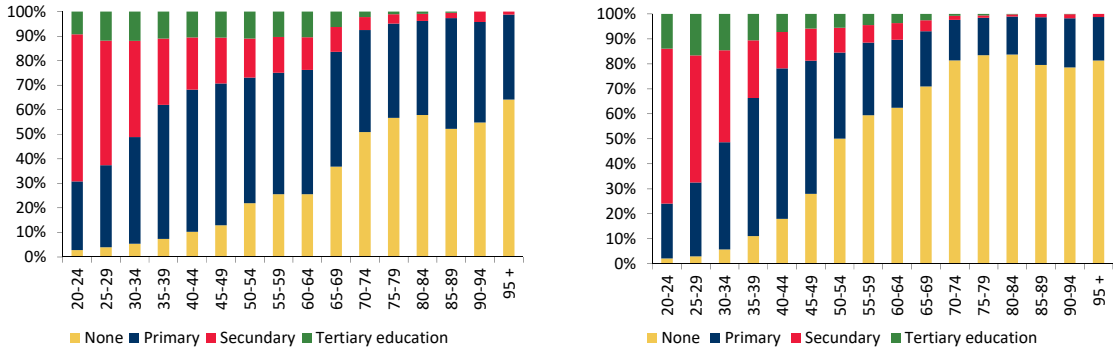

(A) Men (B) Women

FIGURE 12: Distribution of population by educational attainment level in 2010 | In percentage
Source: INE CV (Censos 2010).
Note: Data refers to men and women aged 20 and over.

variables. Therefore it is prudent to use several methods to evaluate the cyclical position of the economy and to evaluate a comprehensive set of economic variables.

Potential output estimates in this article are based on a model of unobserved components in which GDP, the unemployment rate and inflation are decomposed in unobserved trend and cycle components. This model includes a Phillips Curve (relationship between inflation gap and output gap) and an Okun Law (relationship between output gap and unemployment gap). The model was estimated using annual data from 1980 to 2018 and IMF estimate for 2019, assuming three hypotheses for the specification of the unobserved components.

Firstly, it was considered that trend components are not influenced ex-ante by any observed data (designated "Free model"). Secondly, it was considered that the estimates for unobserved variables are partially influenced by low frequency movements, which are based on Hodrick-Prescott (HP filters) with smoothing coefficients equal to 100 ("100 model") and 6.25 ("6.25 model").[21]

The results suggest that the potential output of Cabo Verde's economy presented a deceleration trajectory since late 1990s, recording presently a growth rate around 3.5% (the estimates of the different models point to an interval of growth rates between 2.9% and 4.5%), which compares with an average of 5.3% in 1990 (Figure 13a). It is worth mentioning that in the Free model the estimated trend is close to the one of a HP filter with a smoothing coefficient of 100.

The deceleration of potential output occurred simultaneously with a reduction in trend inflation, measured by the growth rate of the GDP deflator, to values around 1.2% (average of 5.6% in 1990 and of 2.4% in the end of 1990s) (Figure 13b), while the trend unemployment rate reduced from values around 19% to around 13% (Figure 14a). In the case of the unemployment rate trend component, the Free model only suggests

21. UA smoothing coefficient equal to 6.25 is nearly equivalent to a smoothing coefficient of 1,600 when working with quarterly data (see Ravn and Uhlig (2002)). Estimation details can be found in the Appendix.

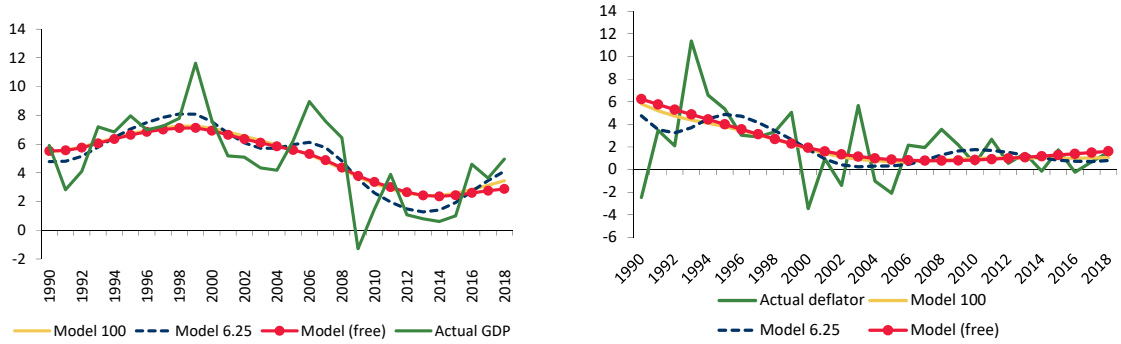

(A) Observed and trend GDP (B) Observed and trend GDP deflator

FIGURE 13: Observed and trend variables | In percentage

Source: BCV and authors' calculations.

Note: Values correspond to annual changes.

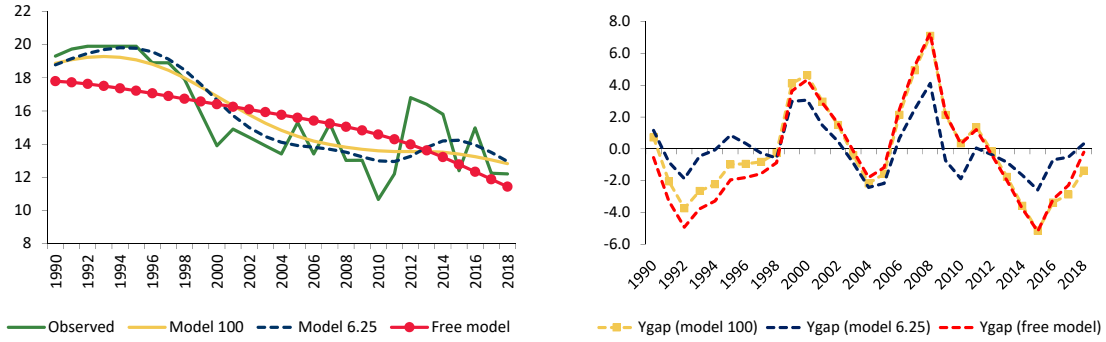

(A) Observed and trend unemployment rate (B) Output gap

FIGURE 14: Labour and product markets

Source: BCV and authors' calculations.

Notes: Unemployment rate is in percentage of active population. Output gap is in percentage of observed GDP.

a decreasing path. In this context, it should be mentioned that the importance of the informal sector in Cabo Verde creates statistical challenges to assess labour market variables and, in particular, historical unemployment rates in distant years. Since 2009, labour market statistics are produced by the National Statistics Institute of Cabo Verde, based on the recommendations of the Resolution of the 13th International Conference of Labour Statistics of the *Bureau Internacional du Travail*, being statistically compatible since 2006.

Output gap estimates (ygap) obtained through the different model parameterizations share some characteristics. Among them, it does stand out that estimates for 2008 systematically correspond to the highest levels (Figure 14b). For the most recent period, estimates suggest that output gaps are closer to zero, after a convergence movement that started in 2015 (from negative output gaps between -2.5% and -5.0%).

It is worth mentioning that these estimates present nevertheless some dispersion, which is a sign of the uncertainty associated to the calculation of these indicators. Between 1990 and 1998, for example, alternative parameterizations give rise to very distinct signs and amplitudes, which highlight the necessity to complement the analysis of the cyclical position of the Cabo Verde's economy with auxiliary information. On the other hand, the output gaps of the Free model and of the 100 model are relatively close, but their relation with the observed unemployment and associated trend values is considerably different (Figure 14a).

Despite the uncertainty, results show decreasing trends in the inflation rate and in the unemployment rate, while potential GDP, although depicting a robust growth in recent years, is below the estimate for the period before the international financial crisis.

In this context, initiatives like the *Plano Estratégico de Desenvolvimento Sustentável* (PEDS) 2017–2021 2021 that Cabo Verdean authorities implemented, aiming to promote economic growth, improve living conditions and public debt sustainability, are particularly relevant.[22] The proposed growth model is based on the promotion of private investment, through the increase of confidence levels and reduction of economic costs, complemented by social and solidarity policies. Among the planned reforms, it is worth mentioning the ones in the financial area, like the reinforcement of supervision, the improvement in the central balance sheet database, credit register and collateral register databases, and also the implementation of recommendations regarding anti money laundering and terrorist financing.

PEDS includes also reforms regarding the improvement of the business environment to support the private sector development, establishing a "financing ecosystem" (which includes guarantees, risk capital and support to project development), and the simplification of procedures to create enterprises. Nevertheless, its timely implementation and close monitoring is crucial to the plan's success.

The approval, in July 2019, of the IMF's Policy Coordination Instrument (PCI), in place for 18 months, i.e. until January 2021, is an important additional tool. The IMF will ensure technical support in order to enhance macroeconomic stability through budgetary consolidation and to promote structural reforms. The budgetary program is based on the improvement of primary balances and ending, in time, with State transfers to public enterprises which run deficits, while a restructuring program takes place, namely through privatizations and licensing of some services. The PCI does not include IMF financial support but, according to the IMF evaluation, the successful conclusion of half-yearly reviews will help signal the commitment of Cabo Verde's authorities in the prosecution of structural reforms.[23]

22. See Ministério das Finanças de Cabo Verde (2018).

23. Recall that the Policy Coordination Instrument (2019) followed several programmes between Cabo Verde and the IMF (with or without financial support: Stand-By Arrangement (1998/2000), Poverty Reduction and Growth Facility (2002/2005) and Policy Support Instrument (2006/2010 and 2010/2012).

6. Final remarks

The goals of the Exchange Rate Cooperation Agreement established between Portugal and Cabo Verde — to promote macroeconomic and financial stability and to increase economic openness — have been achieved with a level of success that justifies in-depth analysis and represents a case study.

The increase in *per capita* income, against a background of higher openness, was one of the most important factors that shaped the last two decades. Among the future challenges, stands out the need to continue the fiscal consolidation process and to enhance economic growth. Efforts that increase labour force qualifications, alongside policies concerned with fostering productive investment, emerge as priorities. The authorities of Cabo Verde should take advantage of ongoing initiatives, such as the *Plano Estratégico de Desenvolvimento Sustentável* and the Policy Coordination Instrument, which aim to promote the population well-being through reforms that adjust the structural features of Cabo Verde to the identified challenges.

Appendix: Methodologies to estimate potential output and trend unemployment rate

The model used in this article decomposes the GDP (y), unemployment rate (u) and inflation (π) in a trend component ($trend$) and a cyclical component (gap), namely $y = y_{trend} + y_{gap}$, $u = u_{trend} + u_{gap}$ e $\pi = \pi_{trend} + \pi_{gap}$.

Trends are described by equations

$$
\begin{aligned}
y_{trend} &= y_{trend(-1)} + I_{\Delta y} && \text{(assuming } I_{\Delta y} = \rho_1 \Delta I_y + (1-\rho_1) I_{\Delta y(-1)} + \varepsilon_y\text{)}, \\
u_{trend} &= u_{trend(-1)} + I_{\Delta u} && (I_{\Delta u} = \rho_2 \Delta I_u + (1-\rho_2) I_{\Delta u(-1)} + \varepsilon_u), \\
\pi_{trend} &= \pi_{trend(-1)} + I_{\Delta \pi} && (I_{\Delta \pi} = \rho_2 \Delta I_\pi + (1-\rho_2) I_{\Delta \pi(-1)} + \varepsilon_\pi),
\end{aligned}
$$

in which the unobserved variables $I_{\Delta y}$, $I_{\Delta u}$ e $I_{\Delta \pi}$ are influenced ex ante by estimates of $\Delta I_y = I_y - I_{y(-1)}$, $\Delta I_u = I_u - I_{u(-1)}$ e $\Delta I_y = \pi - \pi(-1)$, when $\rho_1, \rho_2, \rho_3 \neq 0$, while ε_y, ε_u e ε_π are independent and identically distributed (iid) shocks, with normal distribution and zero mean. Subscript (-1) identifies variables in the previous period. The model was estimate assuming that I_y and I_u correspond to the result of a Hodrick-Prescott filter over GDP and the unemployment rate, respectively, with a smoothing coefficient equal to 100 ("100 model") and to 6.25 ("6.25 model"). The "Free model" assumes that $\rho_1, \rho_2, \rho_3 = 0$.

Cyclical components are connected through a Phillips Curve and an Okun's Law; the output gap is described by an autoregressive process. In particular,

$$
\begin{aligned}
\pi_{gap} &= \beta_1 y_{gap(-1)} + \varepsilon_{\pi_{gap}}, \\
u_{gap} &= \gamma_1 u_{gap(-1)} + \varepsilon_{u_{gap}}, \\
y_{gap} &= \alpha_1 y_{gap(-1)} + \varepsilon_{y_{gap}},
\end{aligned}
$$

where $\varepsilon_{\pi_{gap}}$, $\varepsilon_{u_{gap}}$ e $\varepsilon_{y_{gap}}$ are iid shocks, with normal distribution and zero mean. The models were parameterized with Bayesian techniques, and the results are reported in Table A.1.

	***Prior* distribution**		***Posterior* distribution - median**		
Parameters	Average	Distribution	*Free model*	*6.25 model*	*100 model*
α_1	0.5	β	0.64	0.49	0.59
γ_1	0.5	β	0.52	0.49	0.47
γ_2	0.5	Γ	0.42	0.45	0.41
ρ_1	0.5	β	0	0.75	0.69
ρ_2	0.5	β	0	0.62	0.62
ρ_3	0.5	Γ	0	0.69	0.62
β_1	0.5	Γ	0.28	0.35	0.28

TABLE A.1. Prior and posterior distributions

Source: Authors' calculations.

Notes: Beta and Gama distributions are identified by β and Γ, respectively. The posterior distributions are computed with 1 million draws, from which we discard 40%. Inflation is measured by the growth rate of GDP deflator, nevertheless the results are qualitatively similar when using the CPI. Results are based on IRIS, a Matlab toolbox available at http.//www.iris-toolbox.com.

References

Amador, J., C. Braz, M.M. Campos, S. Sazedj, and L. Wemans (2016). "Public debt sustainability: methodologies and debates in European institutions." Occasional Paper 1, Banco de Portugal.

Correia, Olavo (2002). "Cabo Verde no contexto das relações com a União Europeia e Portugal — ponto de situação e perspetivas." Mimeo (Novembro).

Cruz-Rodriguez, A. (2013). "Choosing and assessing exchange rate regimes: a survey of the literature." *Revista de Análisis Económico*, 28(2), 37–61.

Duarte, C., S. Sazedj, and J.R. Maria (2019). "Trends and cycles under changing economic conditions." Working Paper 7, Banco de Portugal.

Ghosh, A.R. and J.D. Ostry (2009). "Choosing an Exchange Rate Regime." Finance and Development, International Monetary Fund.

Governo de Cabo Verde (1997). Grandes Opções do Plano 1997–2000.

Instituto Nacional de Estatística de Cabo Verde (2014). Estatísticas das Migrações.

Instituto Nacional de Estatística de Cabo Verde (2015). Estatísticas das Migrações.

Ministério das Finanças de Cabo Verde (2018). PEDS - Plano Estratégico de Desenvolvimento Sustentável 2017/2021.

Ravn, Morten O. and Harald Uhlig (2002). "On adjusting the Hodrick-Prescott filter for the frequency of observations." *The Review of Economics and Statistics*, 84(2), 371–380.

Saramago, L. (1998). "Acordo cambial entre Portugal e Cabo Verde: enquadramento, discussão e propostas." Mimeo, Banco de Portugal.

Saramago, L. (2018). "Acordo de Cooperação Cambial entre Portugal e Cabo Verde: 20 anos - contornos, evolução e impacto." , Banco de Portugal.

Saramago, L. and F. Heitor (2015). "Acordos de cooperação cambial e económica na esfera da lusofonia - os casos luso-cabo-verdiano e luso-santomense." Cadernos de Cooperação 6, Banco de Portugal.

Serviço de Estrangeiros e Fronteiras (2019). Relatório de Emigração, Fronteiras e Asilo - 2018, Portugal.

Transparency International (2020). Corruption Perception Index 2019 for Cabo Verde, (available at https://www.transparency.org/).

World Bank (2019). "Worldwide Governance Indicators." Country Data for Cabo Verde 1996-2018, (available at https://info.worldbank.org/governance/wgi/).

Non-technical summary

April 2020

The granularity of Portuguese firm-level exports

Sónia Cabral, Carlos Melo Gouveia, Cristina Manteu

Recent decades have been characterised by strong growth of international trade worldwide and Portugal is no exception. In the last 25 years, the total value of Portuguese exports of goods more than quadrupled and the number of exporting firms more than doubled, from around 12,200 in 1993 to around 26,000 in 2017. In fact, the positive export performance of the Portuguese economy has been a key variable in the latest economic recovery. Hence, understanding the dynamics of exports is fundamental to predict economic growth and better inform policy-decisions in this area.

In spite of the strong rise in the number of exporting firms, the values of Portuguese exports are very concentrated in a few large firms. Evidence for other countries shows that, if markets are dominated by a small number of large firms, specific events that affect these major firms can impact the evolution of economic aggregates. A statistical procedure to analyse a distribution where activity is extremely concentrated in a few top firms is to estimate a power law or Pareto distribution, with lower exponents associated with higher concentration.

Our article aims at answering two questions using Portuguese firm-level exports of goods from 1993 to 2017. First, we test if Portuguese exports are so concentrated in the major firms that the tail of the distribution can be described by a power law with a very low exponent. Second, we assess how the idiosyncratic movements of the largest exporters affect the growth rate of aggregate exports.

Starting with the first question, we apply three different methods of estimation of the Pareto exponent of the tail of the distribution of Portuguese firm-level exports. Irrespective of the method, the estimated exponents are very small, pointing to an extreme concentration of export values in a few firms.

Regarding the second question, the firm-specific shocks to exports are defined as the difference between the individual firm's growth rate of exports and the average rate of a small group of major firms. The granular residual is the weighted average of these firm-specific shocks, as in Gabaix (2011). The contribution of the granular residual to aggregate fluctuations is obtained by regressing the rate of change of total exports on the granular shock and examining its explanatory power. We find that idiosyncratic shocks to the largest firms are relevant for total Portuguese export growth over time, accounting for more than one-third of aggregate fluctuations.

This means that volatility at the firm-level can affect Portuguese aggregate export dynamics, as illustrated for the case of the top 1000 exporters in the figure below.

This result has implications for the monitoring and forecasting of Portuguese exports. Understanding the evolution of large firms and the specific shocks they are subject to can shed light on the behaviour of aggregate exports over time and help forecast their future growth.

FIGURE 1: Granular residual of the top 1000 exporting firms and aggregate growth of Portuguese exports (1994-2017)

Note: The line results from regressing the rate of change of total Portuguese exports from 1994 to 2017 on the granular residual of the top 1000 exporting firms.

The granularity of Portuguese firm-level exports

Sónia Cabral
Banco de Portugal

Carlos Melo Gouveia
Banco de Portugal

Cristina Manteu
Banco de Portugal

April 2020

Abstract
International trade flows are highly concentrated in the top units of analysis. In this paper, we study the firm-size distribution of Portuguese exports of goods over the period 1993-2017 and examine its impact on aggregate volatility. We fit power law relations using three different methods. The estimated Pareto exponents of Portuguese firm-level exports are very small, pointing to an extreme concentration of export values in the largest firms. Moreover, the Pareto estimates tend to decrease over time. We then investigate if the granular behaviour of aggregate outcomes of Gabaix (2011) is present in Portuguese firm-level exports. The results show that the idiosyncratic shocks to the largest firms may have significant effects on the dynamics of total exports. (JEL: F10, F14)

1. Introduction

Changes in total exports are usually explained by aggregate shocks to external demand and competitiveness, but often the understanding of year-to-year movements of aggregate exports requires examining individual firm behaviour. The fact that a few large firms account for a disproportionate share of activity has implications on the interpretation of several economic phenomena and international trade is no exception. Trade flows are highly concentrated in a few top firms. Such a firm-size distribution is well represented by a power law, with lower exponents associated with higher concentration and fatter tails. The seminal paper of Gabaix (2011) shows that, if the distribution of firm size is very fat-tailed, idiosyncratic shocks to a few very large firms can explain an important part of aggregate fluctuations.

Portugal is a small and open economy where large firms account for a substantial share of total exports and, hence, specific shocks to these firms may have a strong impact

Acknowledgements: The authors are especially grateful to Paulo M. M. Rodrigues for his comments and invaluable help regarding the Nicolau and Rodrigues (2019) estimator. We also thank the editor, Pedro Duarte Neves, and Nuno Alves, João Amador, António Antunes, Cláudia Duarte and participants in an internal seminar of the Economics and Research Department of Banco de Portugal for their useful comments and suggestions. The analyses, opinions, and findings expressed in this article are those of the authors and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.
Note: This article was submitted for approval on February 21, 2020, before the World Health Organization declared the COVID-19 a pandemic on March 11.

E-mail: scabral@bportugal.pt; cgouveia@bportugal.pt; mcvmanteu@bportugal.pt

in the aggregate. There is plenty of anecdotal evidence on this type of effects in Portugal. For instance, a temporary interruption of activities in a refinery firm led to a significant decline in the volume of fuel exports in 2014, while an increase in the productive capacity of an automotive plant had a substantial contribution to the gains of Portuguese export market share in 2018.

This article analyses the granular hypothesis for Portuguese aggregate exports of goods using micro data. Can the tail of the distribution of Portuguese firm-level exports be modelled by a power law with a sufficiently low exponent? If so, do the idiosyncratic movements of the largest exporters explain a significant part of the growth rate of aggregate exports? The answer to both questions is yes, as we detail below.

We use three different methods of estimation of the Pareto exponent of the tail of the distribution of Portuguese firm-level exports in each year from 1993 to 2017. Two are the most popular methods in the literature – the conditional maximum likelihood estimator of Hill (1975) and the log-rank, log-size OLS regression with the optimal shift of Gabaix and Ibragimov (2011) – and the third is the recent estimator proposed by Nicolau and Rodrigues (2019). Irrespective of the method, the estimated Pareto exponents are very small, in some cases close to 1, pointing to an extreme concentration of export values in a few firms. Moreover, the estimates of the Pareto exponent tend to decline over time, in particular until 2009, which indicates an increase in the concentration of the values of exports in the largest firms.

Having established the significant role of top firms, we then examine how this concentration of exports in a few very large firms has an impact in the aggregate growth of Portuguese exports. Following the empirical approach of the granular residual of Gabaix (2011), we find that the idiosyncratic shocks to the largest firms have significant effects on the dynamics of total exports, explaining more than one-third of aggregate fluctuations.

The article is organised as follows. Section 2 discusses some of the related literature that frames this study. Section 3 describes the database and presents some descriptive statistics. Section 4 estimates the Pareto exponent of the tail of the distribution of Portuguese firm-level exports over time. After a short presentation of the methodology, Section 5 investigates the role of idiosyncratic shocks to the largest firms for the aggregate change in Portuguese exports. Finally, Section 6 presents some concluding remarks.

2. Related literature

The review of the vast empirical literature on power laws in economics and their implications for the behaviour of many economic and financial variables is beyond the scope of this article. In fact, there are already several extensive reviews of the literature on these issues that we mention below. Instead, this section offers a non-exhaustive list of references in different strands of the literature that are related to our study and provide a framework for our analysis.

There is a large body of literature in economics that deals with power law distributions. Gabaix (2009, 2016) provide two very intelligible surveys on empirical power laws in economics, discussing their theoretical mechanisms and their utility for the understanding of several economic phenomena, including the granular origins of aggregate fluctuations. Newman (2005) provides a more general discussion of how power law distributions occur in an extraordinarily diverse range of areas from physics, biology, computer science to economics and finance, demography and other social sciences. The great interest in these type of distributions has also led to the refinement and development of a number of tail index estimators over the years (see, for instance, Embrechts *et al.* (2012) for an extensive review of these methods).

The important role of the top units in a distribution links directly with the granular hypothesis of aggregate fluctuations of Gabaix (2011), which provides a micro-foundation for aggregate shocks. The basic idea is that idiosyncratic shocks to firms can generate aggregate fluctuations, i.e., specific events that affect the largest firms can impact the evolution of the economic aggregate of which they are part. The key factor here is the high concentration of economic activity in the top "grains" (firms, industries, or products), with size distributions that can be well fitted by power laws. The "averaging out" argument of firm-level shocks breaks down if the size distribution of firms is very fat-tailed. In this case, random shocks to the largest firms can affect total output in a noticeable way. Gabaix (2011) shows that idiosyncratic movements of the top 100 US firms explain about one-third of variations in output growth.

After the seminal paper of Gabaix (2011), there is a growing literature that looks at firms to understand aggregate outcomes. For instance, using the empirical approach of Gabaix (2011), granular firms are found to account for a significant part of business cycle fluctuations in Germany (Wagner 2012), Spain (Blanco-Arroyo and Alfarano 2017 and Blanco-Arroyo *et al.* 2018), Italy (Gnocato and Rondinelli 2018), Canada (Karasik *et al.* 2016), Australia (Miranda-Pinto and Shen 2019) and eight European countries (Ebeke and Eklou 2017). Using detailed data for sales of French firms to different markets, di Giovanni *et al.* (2014) find that firm-specific shocks contribute more to aggregate volatility than sector-destination shocks do. Friberg and Sanctuary (2016) replicate this analysis for Sweden and also find that the firm-specific component contributes substantially to the growth of total sales and exports.

Other papers have built on the granular hypothesis and developed models in which aggregate fluctuations arise from shocks to individual firms, because the firm-size distribution is extremely fat-tailed. di Giovanni and Levchenko (2012) extend the granular approach to firms in international trade and propose a new channel through which international trade affects aggregate volatility. They show that the preponderance of large firms and their role in aggregate volatility can help explain two empirical regularities: smaller countries are more volatile; and more open countries have higher volatility. Opening to trade increases the importance of large firms, thus raising macroeconomic volatility. More recently, di Giovanni *et al.* (2017, 2018) show that idiosyncratic shocks to the granular firms have an impact on international business cycle comovement; Gaubert and Itskhoki (2018) propose and quantify a granular multi-sector model of trade and show that idiosyncratic firm dynamics account for most of the

evolution of a country's comparative advantage over time; Carvalho and Grassi (2019) extend the work of Gabaix (2011) and build a quantitative firm dynamics model where aggregate fluctuations are caused by firm-level disturbances alone.

Current international trade literature has already established the dominance of large firms in aggregate exports. Empirical evidence on a highly skewed exporter-size distribution is provided, for instance, by Mayer and Ottaviano (2008) for seven European countries, Bernard *et al.* (2009) for the US, Eaton *et al.* (2011) for France, and Freund and Pierola (2015) for 32 developing countries. Our article is mostly related to Wagner (2013) that uses data on total exports at the firm-level and finds that idiosyncratic shocks to very large German manufacturing firms played a decisive role in the export collapse of 2008-2009. Our article is also connected to del Rosal (2013, 2018). We fit a power law distribution to exports as in del Rosal (2018) and estimate the granular residual using the approach of Gabaix (2011) as in del Rosal (2013), but his analyses are at product-level for several European countries and ours is at the firm-level for Portugal.

Finally, our article contributes to the empirical literature on Portuguese international trade in goods using micro-level data. Some examples of studies using the same database of this article include Amador and Opromolla (2013, 2017) who analyse the intensive and extensive margins of Portuguese exports in the firm, destination and product dimensions; Nagengast (2019) who assesses the importance of product and destination shocks to the variation of total exports; Bastos and Silva (2012) who find a positive effect of migrant networks on export participation and intensity; Mion and Opromolla (2014) who show that the export experience gained by managers in previous firms has a positive impact on the export performance of their current firm; Bastos and Silva (2010) who find that more productive firms sell larger volumes at higher prices to a given destination; Bastos *et al.* (2018b) who show that exporting to richer countries leads firms to pay higher prices for inputs; Esteves *et al.* (2018) who estimate a negative relation between domestic demand and firms' exports; Bastos *et al.* (2018a) who document new facts on the joint evolution of firm performance and prices over the life cycle.

3. Data and exploratory analysis

Statistics Portugal (*Instituto Nacional de Estatística*, Portuguese acronym: INE) collects, on a monthly basis, data on export and import transactions of goods by firms that are located in Portugal to compute the official international trade statistics. In this article, we use the total values of exports reported by firms from 1993 to 2017 aggregated at the annual level and expressed in current euros. To reach the total official value of exports published, the database comprises some estimates done by INE for non-reporting firms, which are included in fictional identifiers. We exclude this information from the analysis and, hence, use only the export values reported by firms. The firm-level data used in this article covers around 97 percent of total exports as published in the official Portuguese statistics of international trade.

Figure 1 illustrates the strong growth of the Portuguese exporting sector over this period: the total value exported more than quadrupled, from about 12,300 million euros

in 1993 to around 53,700 million euros in 2017, and the number of exporting firms more than doubled, from around 12,200 in 1993 to around 26,000 in 2017. Another feature that stands out in Figure 1 is that, as found for other countries, the great trade collapse of 2009 resulted mostly from the decline in exports of firms that continued to export (i.e., the so-called intensive margin) and not from the reduction of the number of exporting firms.[1]

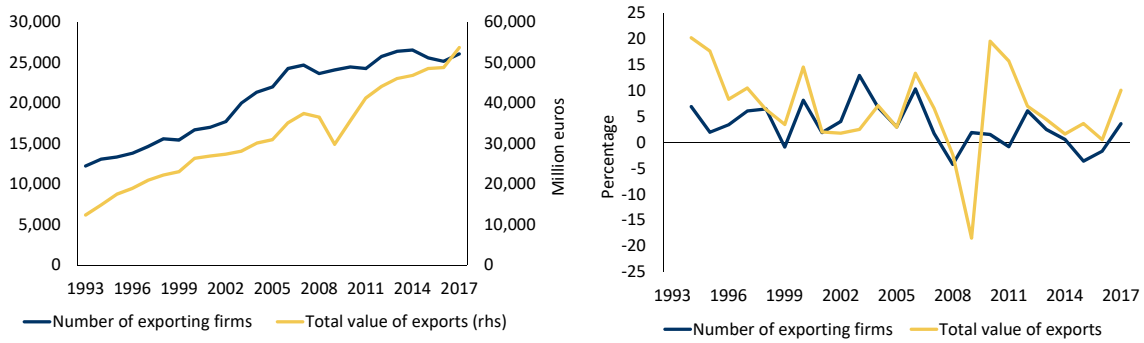

(A) Levels (B) Rates of change

FIGURE 1: Portuguese exports - number of exporting firms and total value (1993-2017)

Are the values of Portuguese exports concentrated among firms? Has this feature changed over time? In the next section, we try to answer these questions by estimating the shape parameter of the tail of the distribution of firm-level exports, but some preliminary evidence is presented in Figure 2. The simplest indicators of export concentration are the shares of the major exporting firms depicted in Panel a). These shares illustrate the high concentration of Portuguese exports in a small group of firms. For instance, the top 250 exporters in 1993 (2 percent of firms) account for around half of total exports, while the main 1000 exporters (8 percent of firms) represent more than 75 percent of the total export value in 1993. In spite of the strong rise in the number of exporting firms, the shares of top firms in the total value of exports have also increased, especially until 2009. In 2017, the major 250 and 1000 exporting firms, which represent only around 1 and 4 percent of exporters, account for almost 60 and 80 percent of total exports, respectively.

Complementarily, Panel b) of Figure 2 reports some traditional indices of inequality for the whole sample, namely the Gini coefficient, the dissimilarity index, and the Theil T_T entropy index, which can also be used to assess export concentration.[2] All three indices have their minimum values at zero (if all firms exported the same value) and

1. For evidence on the importance of adjustments at the intensive margin during the 2008–2009 trade collapse, see, for instance, Bricongne *et al.* (2012) for France, Wagner (2013) for Germany, Behrens *et al.* (2013) for Belgium, Eppinger *et al.* (2018) for Spain, and Békés *et al.* (2011) for seven European countries.

2. Denote the share of firm i in total exports as p_i and N as the total number of firms, the Theil index is $T_T = lnN - \sum_{i=1}^{N} p_i * ln(1/p_i)$ and the dissimilarity index is $D = 1/2 * \sum_{i=1}^{N} |p_i - 1/N|$. After sorting firms in ascending order of their export value x_i, the Gini index is $G = \frac{2\sum_{i=1}^{N} i*x_i}{N\sum_{i=1}^{N} x_i} - \frac{N+1}{N}$.

increase as the values of exports become more concentrated in some firms. Again, the evolution of these indicators points to an increase in the concentration of export values among firms over time.

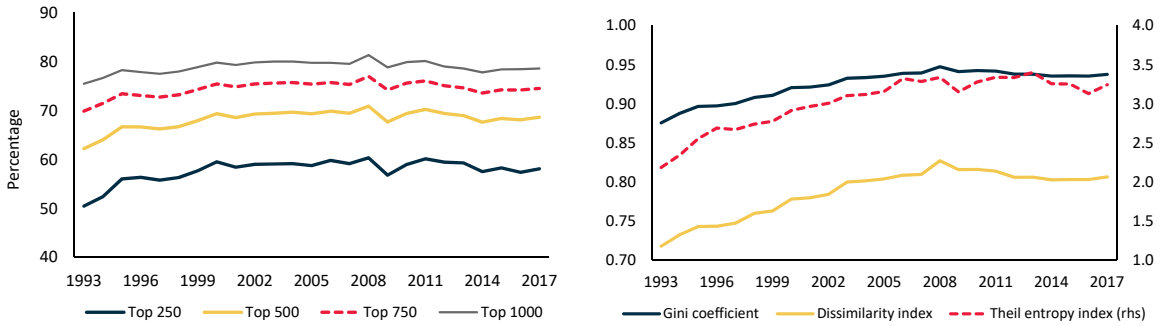

(A) Shares of top firms in total exports (B) Inequality measures

FIGURE 2: Concentration of Portuguese exports (1993-2017)

We finish this section with some additional information on the characteristics of the top Portuguese exporters. More precisely, we examine their frequency and export representativity by sector of economic activity, defined using the Portuguese industrial classification – *Classificação Portuguesa das Actividades Económicas* (CAE) Rev 3 at the 2-digits level. Table 1 illustrates the sectoral distribution of the top 250 exporters, in terms of both number of firms and value of exports over the 1993-2017 period. To facilitate the presentation, only the main fifteen sectors of activity are included in the Table 1 and sorted in descending order of their percentage shares.

Two sectors stand out in Table 1: motor vehicles and wholesale trade.[3] Each of these sectors accounts for around 11 percent of the number of largest firms in this period. The relevance of the sector of motor vehicles in terms of the value of exports of the top firms is higher (20.7 percent), while the shares of wholesalers are similar in terms of value of exports and number of firms. The importance of wholesalers in Portuguese top exporters is in line with evidence reported for other countries and also on previous studies on Portuguese international trade: many firms classified in service sectors operate and play a significant role in exports of goods.[4]

Fourteen of the fifteen major sectors of Table 1 are common in terms of their shares in the number of largest exporters and in the values of exports of these top firms. The two exceptions are footwear and leather, and coke, refined petroleum and fuels. The latter sector, in particular, represents almost 9 percent of the export value of the largest firms but less than 0.5 percent in terms of the number of top firms, pointing to the existence of a single major exporter in this sector.

3. According to CAE Rev 3, wholesale trade comprises the activity of resale (without transformation) of goods to traders (retailers or other wholesalers), to industrial, institutional and professional users, to intermediaries, but not to final consumers. The goods can be resold as they were acquired, or after carrying out some operations associated with wholesale trade, such as sorting in lots, packaging and bottling.

4. See, for instance, Crozet *et al.* (2013) for the specific role of wholesalers in exports of goods and Banco de Portugal (2016) for evidence on Portuguese service firms participating in international trade of goods.

Shares in the number of largest exporters			Shares in the value of exports of the largest firms		
Code		Share	Code		Share
29	Motor vehicles	11.1	29	Motor vehicles	20.7
46	Wholesale trade	11.0	46	Wholesale trade	10.5
13	Textiles	8.2	19	Coke, refined petroleum and fuels	8.9
14	Apparel	6.3	26	Computer and electronics	6.7
10	Food products	5.1	17	Paper and paper products	5.6
23	Other non-metallic mineral products	5.0	20	Chemicals, except pharmaceuticals	5.1
16	Wood and cork	4.9	27	Electrical equipment	4.2
27	Electrical equipment	4.6	22	Rubber and plastic products	4.1
20	Chemicals, except pharmaceuticals	4.5	13	Textiles	3.7
22	Rubber and plastic products	4.1	24	Basic metals	3.6
15	Footwear and leather	3.8	16	Wood and cork	3.0
25	Metal products, except machinery	3.7	10	Food products	2.9
17	Paper and paper products	3.1	23	Other non-metallic mineral products	2.6
24	Basic metals	3.1	14	Apparel	2.1
26	Computer and electronics	3.0	25	Metal products, except machinery	1.9
	Other sectors	18.5		Other sectors	14.5
	Total	100		Total	100

TABLE 1. Distribution of the largest 250 Portuguese exporters by sector of activity - percentage shares in the number and exports of the largest firms, 1993-2017

Notes: The sectors are defined at the 2-digits level of CAE Rev3. Only the fifteen sectors with higher shares are included and sorted in descending order of their shares. The percentage shares are computed using all 25 years from 1993 to 2017 and the largest 250 Portuguese exporting firms in each year.

4. The power-law of Portuguese firm-level exports

"When the probability of measuring a particular value of some quantity varies inversely as a power of that value, the quantity is said to follow a power law, also known as Pareto distribution." (Newman 2005, page 1).

This section begins by describing how the tail of the distribution of firm-level exports can be approximated by a power law relation. Figure 3 presents the exports of the Portuguese top 1000 firms in 1993 and 2017. Both axis are in logarithmic scales. The x-axis displays the firm rank, i.e., the first one is the firm that exported the most in the reference year and so on, whereas in the y-axis, its exports in million euros are depicted. The fact that, in both years, the relationship can be represented closely by a straight-line (R-squared in 1993 = 0.991 and R-squared in 2017 = 0.997) suggests that firm exports in Portugal follow a power law.

Analytically, let X_{it} be the export value of firm i in a given year t and $P(X_{it} > x)$ the probability that X_{it} takes a value equal to or greater than x. This variable follows a power law or Pareto distribution if the counter cumulative distribution function or empirical survival function is:

$$\overline{F}(x) = 1 - F(x) = P(X_{it} > x) = \left(\frac{x_0}{x}\right)^{\alpha}, \quad \text{with } x > x_0 > 0, \alpha > 0, \tag{1}$$

with x_0 denoting the lower bound to the power law behaviour of the distribution. The exponent α, also known as the Pareto or tail exponent, is the key parameter here. It provides a measure of how heavy the tails are, i.e., how concentrated the variable is in

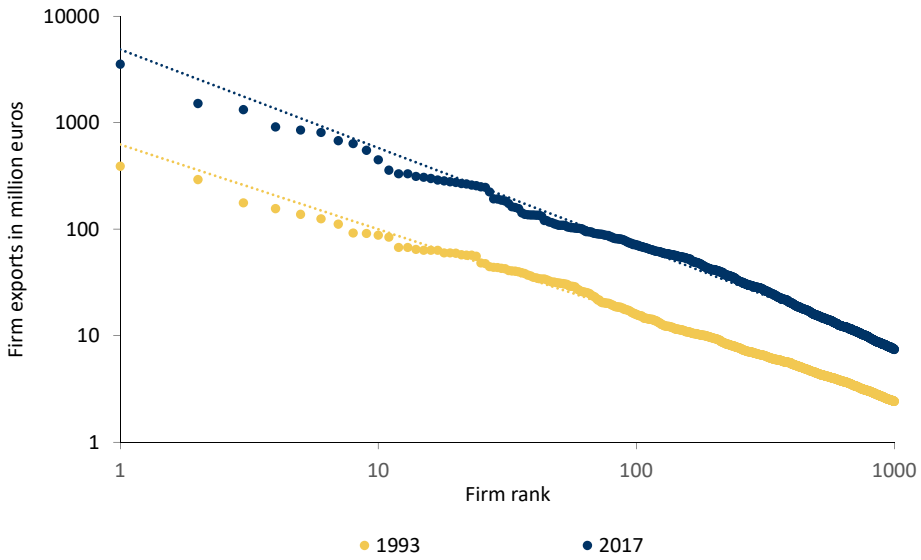

FIGURE 3: Portuguese exports of the top 1000 firms in 1993 and 2017

Note: The dots are the empirical data plotted on logarithmic scales. The lines illustrate a power law fit regressing firm exports on firm rank.

the top units. The lower the exponent α, the fatter the tails of the distribution, the greater the degree of inequality in the distribution and the higher the probability of finding very high values. For instance, a Zipf's law states that $\alpha \approx 1$.[5]

Given the number N of units in the upper tail implied by the value of x_0, we use three different methods of estimation of α in each year t from 1993 to 2017. Two are the most popular methods in the literature and the other is the recent estimator proposed by Nicolau and Rodrigues (2019).

The first method is the conditional maximum likelihood estimator (MLE) proposed by Hill (1975),

$$\hat{\alpha}_{MLE} = N \left(\sum_{i=1}^{N} lnX_{it} - lnx_0 \right)^{-1}. \tag{2}$$

The second method is the log-rank, log-size OLS regression with the shift proposed by Gabaix and Ibragimov (2011) ($\gamma = 1/2$), where α is the slope of the following regression of the log rank of firm i in year t (r_{it}) on its log size:

$$ln(r_{it} - 1/2) = \beta_0 - \alpha lnX_{it} + \varepsilon_i. \tag{3}$$

5. Zipf's law was originally formulated in terms of quantitative linguistics and is named after the American linguist George Kingsley Zipf, who popularised it. The law states that the frequency of any word in any text is inversely proportional to its rank in the frequency table, i.e., the most frequent word will occur twice as often as the second most frequent word, three times as often as the third most frequent word, etc.. This regularity exists in all languages and, for instance, the distribution of the words used in this article follows a Zipf's law.

Even if a popular way of estimating the Pareto exponent is to run a standard OLS log-log rank-size regression with $\gamma = 0$, this procedure is strongly biased in small samples. Gabaix and Ibragimov (2011) show that applying a shift $\gamma = 1/2$ is optimal in reducing the small sample bias. For future reference, we name the resulting OLS estimate of α of eq. (3) as $\hat{\alpha}_{OLS1/2}$.

A caveat of these two estimation procedures is that they assume that the observations are independent, which is not necessarily the case in most microeconomic data sets. However, there is no consensus in the literature on how to solve this issue. In practice, most applied studies present one of the former estimators, together with a warning that the observations are not necessarily independent and, hence, the standard errors probably underestimate the true standard errors (see Gabaix 2009).

As defined in Equation (1), the Pareto distribution only applies for values equal to or greater than some minimum of value x_0, implying a cut-off still in the upper tail. However, how to choose this cut-off value is also not consensual in the literature. Given that the Hill estimator is especially sensitive to the sample size, Clauset *et al.* (2009) advocate a method for estimating x_0 that uses the Kolmogorov-Smirnov test. However, in practice, most applied studies still rely on a visual goodness of fit or use a simple rule, like the major 1000 units or the top 5 percent of the distribution.

The third method that we apply in this article is the $\hat{\alpha}_{Pareto}$, a new regression-based estimator recently proposed by Nicolau and Rodrigues (2019), which minimises the two caveats mentioned above. This estimator is more resilient to the choice of the sub-sample of large observations used to estimate the Pareto exponent and it performs well under dependence of unknown form. Moreover, the $\hat{\alpha}_{Pareto}$ estimator also provides a bias reduction when compared to the other approaches.

Figure 4 presents our estimation results of the Pareto exponent of the tail of the distribution of Portuguese firm-level exports from 1993 to 2017. Besides using three distinct methods, we also estimate the Pareto exponent for different tail truncation levels. Given that the sample size used in estimation should represent a significant share of the total value of exports, we use the top firms depicted in Panel a) of Figure 2 of the previous section: the main 250, 500, 750 and 1000 exporting firms in each year. An alternative strategy, which we also tested, is to choose a specific upper percentile of the annual distribution of firm-level exports (e.g., the top 5 percent exporters). In our case, the main results are similar, but, to maintain the consistency throughout the article, we chose to present the results based on a fixed number of major exporters. Moreover, we also applied the methodology of Clauset *et al.* (2009) for estimating the lower bound of power-law behaviour in each year. The implied number of firms varies strongly (from a minimum of 276 to a maximum of 1244) but the estimates of the Pareto exponents are consistent with those depicted in Figure 4.[6]

The estimated Pareto exponents of Portuguese firm-level exports are very small, always below 1.4 and in some cases very close to 1, pointing to an extreme concentration of export values in a few firms. These results are in line with evidence on firm-size

6. All estimates are available from the authors upon request.

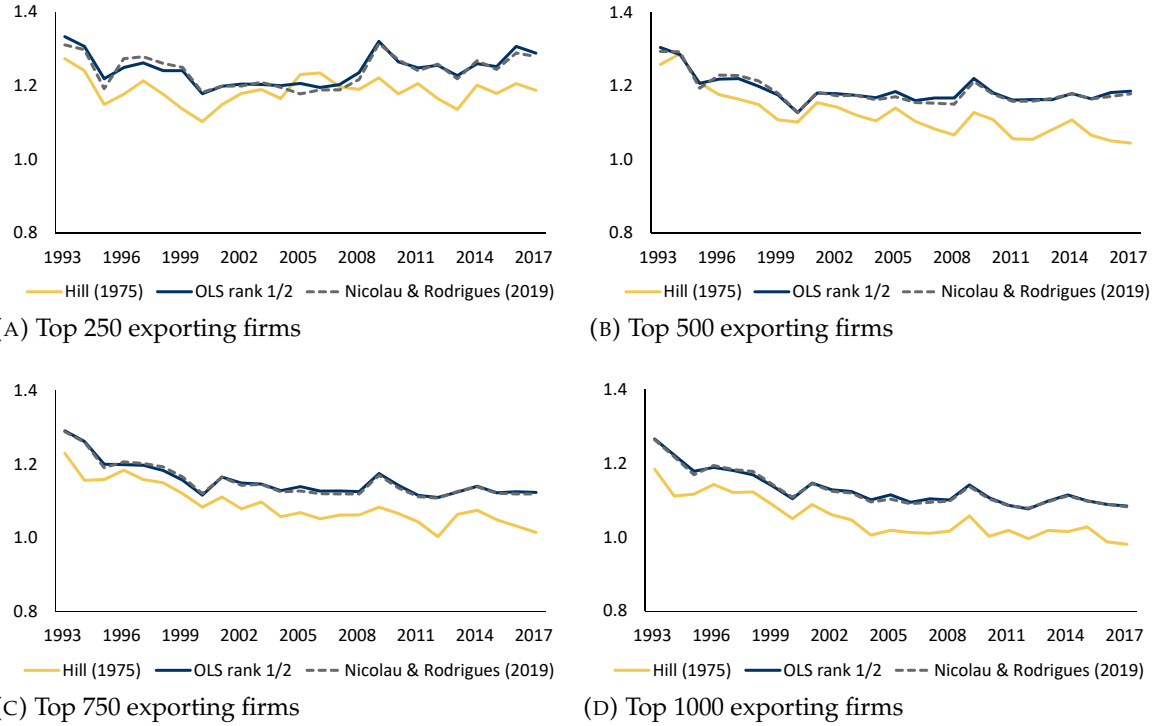

(A) Top 250 exporting firms

(B) Top 500 exporting firms

(C) Top 750 exporting firms

(D) Top 1000 exporting firms

FIGURE 4: Estimates of the Pareto exponent of the tail of the distribution of Portuguese firm-level exports (1993-2017)

Notes: Hill (1975) is the maximum likelihood estimator defined in Equation (2), OLS rank 1/2 is the OLS estimate of the log-rank, log-size regression defined in Equation (3). Four distinct tail thresholds are used in each year: top 250, 500, 750 and 1000 exporting firms.

distributions of other countries and using other proxies for firm-size, such as total sales, number of employees, assets, or revenues. Estimates of the Pareto exponent consistent with a Zipf's law are found, for instance, by Okuyama *et al.* (1999) for Japan, Axtell (2001) for the US, Fujiwara *et al.* (2004) for several European countries, Zhang *et al.* (2009) and Gao *et al.* (2015) for China, di Giovanni *et al.* (2011) for France, and da Silva *et al.* (2018) for Brazil. Our results are also consistent with those of del Rosal (2018). He studies the size distribution of exports at the product-level for the 28 European Union countries and concludes that, given the high concentration of values, power law relations with very low exponents are a good approximation to the data. For instance, considering the top 5 percent products in 2014, the Pareto exponent estimate for Portugal is 1.2, which is in the middle of the range of the estimates obtained by del Rosal (2018) for the several countries.

Another feature evident in Figure 4, in particular until 2009, is that most estimates of the Pareto exponent tend to decline over time. As discussed in Gao *et al.* (2015) that found similar results for Chinese firms using sales and equities from 2001 to 2013, an evolving Pareto coefficient implies that the relative firm sizes are changing. For large firms above the lower bound x_0, the smaller the coefficient, the greater the degree of firm-size inequality and the less homogeneous the relative firm sizes.

The downward path of Figure 4, especially until 2009, is also consistent with the rise of the measures of inequality reported in the previous section, which indicate an increase

in the concentration of the values of exports, i.e., an increase in inequality in the whole sample of Portuguese exporting firms. Note that the Pareto coefficient is an inequality indicator only for the largest exporting firms included in the estimation of the power-law distribution. However, given the high concentration of Portuguese exports in these firms and the tail truncation levels used in the estimation, they account always for more than half of the total values exported.

The fact that a few large firms account for a disproportionate share of activity has implications on the interpretation of several economic phenomena, including aggregate fluctuations as discussed in Section 2. The granular hypothesis of Gabaix (2011) states that, if the distribution of firm size is very fat-tailed, then idiosyncratic shocks to large firms do not cancel out and have an impact on aggregate volatility. This is the subject of the next section of this article. In a different vein, di Giovanni and Levchenko (2013) show that the distribution of firm size is important for evaluating the gains from reductions in entry barriers and trade costs and the relative importance of intensive and extensive margins. Using firm-level data on exports from 32 developing countries, Freund and Pierola (2015) provide evidence on the very high concentration of exports in a few very large firms and on their importance in shaping sectoral trade structures and comparative advantage. This implies that small policy interventions can have big aggregate effects if they change the behaviour of the largest firms.

5. The granularity of Portuguese aggregate exports

In this section, we investigate if the extreme concentration of Portuguese exports in a few very large firms has an impact in the aggregate growth of exports. Are the idiosyncratic shocks to these large firms relevant for total export growth over time? We use the empirical strategy of the granular residual of Gabaix (2011) to try to answer this question.

Let X_{it} denote the export value of firm i in a given year t, $X_t = \sum_i X_{it}$ the value of total Portuguese exports in year t, g_{it} and g_t the respective growth rates, and K the number of granular firms. The growth rate of a firm's exports g_{it} comprises two components: one common to all firms (a macro shock) and one specific to the firm. There are several possible ways to quantify the macro shock. Gabaix (2011) uses a very simple way: the equal-weighted average growth rate of a small subset Q of very large firms. The firm-specific shock is then the portion of the growth rate g_{it} that is unaccounted for by the common shock. Finally, the granular residual Γ_t is defined as the sum of firm-specific shocks in year t, weighted by their size in the previous year. Analytically,

$$\Gamma_t = \sum_{i=1}^{K} \frac{X_{i,t-1}}{X_{t-1}} \hat{\varepsilon}_{it} = \sum_{i=1}^{K} \frac{X_{i,t-1}}{X_{t-1}} (g_{it} - \overline{g_t}) \quad \text{with} \quad \overline{g_t} = Q^{-1} \sum_{i=1}^{Q} g_{it}. \tag{4}$$

Note that, when computing the granular residual Γ_t, we are focusing only on the intensive margin of export growth, i.e., on those firms that export in t and $t-1$ so that a growth rate of their exports can be computed. In addition, the K and Q top firms considered are the largest in terms of their exported value in $t-1$. Moreover, as standard

in this literature, we winsorised the top and bottom 5 percentiles of export growth rates at the firm-level, as an outlier treatment. However, the main results are very similar if we do not winsorise at all or if we use other thresholds of winsorisation, like the top and bottom 1 or 10 percentiles.

After computing the granular residual, the question of interest is whether these idiosyncratic shocks to large exporting firms can impact the growth rate of total exports. The explanatory power of the granular residual Γ_t on aggregate export growth is captured by running the following simple regression:

$$g_t = \beta_0 + \beta_1 \Gamma_t + u_t, \tag{5}$$

where β_0 and β_1 are parameters to be estimated and u_t is an error term. An extended version of the previous equation considers also the lagged values of the granular residual. The coefficient of determination, R-squared, of Equation (5) represents the proportion of the variance for the aggregate annual growth of Portuguese exports that is explained by the granular shocks.

Table 2 presents the regression results with and without a lag of the granular residual for four distinct values of K largest firms with $Q = K$. We tested several numbers for top exporting firms K and $Q \geq K$ and our results are robust to these checks. Moreover, the results are basically unchanged if the dependent variable of Equation (5), the aggregate rate of growth of Portuguese exports g_t, is computed using only the export values of continuing firms.[7]

The idiosyncratic movements of the top firms account for a large fraction of aggregate export fluctuations. If only common shocks were important for the growth rate of total exports, then the R-squared of the regressions in Table 2 would be zero, but it is not. Considering the top 250 firms, the adjusted R-squared shows that the granular residual without any lag can explain, in a statistical sense, around 29 percent of the variability of export growth and up to 40 percent with one lag.[8] The adjusted R-squared increases as the number of top firms rises, reaching around 60 percent with the top 750 and 1000 firms. This is a relatively high value in this literature, even using slightly different methodologies and distinct variables, but it is in line with the cross-country evidence of del Rosal (2013). Using exports at the product-level for several European countries, del Rosal (2013) finds that Portugal is the country with the highest explanatory power of the granular residual, with an adjusted R-squared of more than 30 percent in the regressions with one lag.

Figure 5 illustrates the goodness of fit of Equation (5) by plotting the observed rate of change of total Portuguese exports and the respective fitted values. The granular residuals of the top firms capture very well the variation in the growth of total exports in this period, suggesting that monitoring the qualitative and quantitative data on a panel of major exporters can help to predict aggregate growth.

7. All results are available from the authors upon request.

8. The inclusion of additional lags does not improve the results of Table 2.

	$K=Q=250$		$K=Q=500$		$K=Q=750$		$K=Q=1000$	
Γ_t	2.433**	2.652**	2.641***	2.776***	2.757***	2.751***	2.401***	2.361***
	(0.752)	(0.698)	(0.598)	(0.551)	(0.511)	(0.479)	(0.429)	(0.409)
Γ_{t-1}		1.675*		0.960		0.613		0.643
		(0.687)		(0.554)		(0.476)		(0.402)
Intercept	6.876***	6.746***	6.574***	6.108***	5.971***	5.372***	6.203***	5.727***
	(1.419)	(1.286)	(1.253)	(1.143)	(1.137)	(1.072)	(1.108)	(1.054)
Observations	24	23	24	23	24	23	24	23
R^2	0.322	0.456	0.470	0.562	0.569	0.623	0.588	0.631
Adjusted R^2	0.292	0.402	0.446	0.518	0.549	0.586	0.569	0.594

TABLE 2. Granular residual and aggregate growth of Portuguese exports (1994-2017)

Notes: The table reports the estimation results of Equation (5) where the aggregate growth of Portuguese exports from 1994 to 2017 was regressed on the granular residual Γ_t of four different groups of top firms $K = Q = \{250, 500, 750, 1000\}$. The firms are the largest by their exports in the previous year. Standard errors are in parenthesis. Stars indicate significance levels of 5% (*), 1% (**), and 0.1%(***).

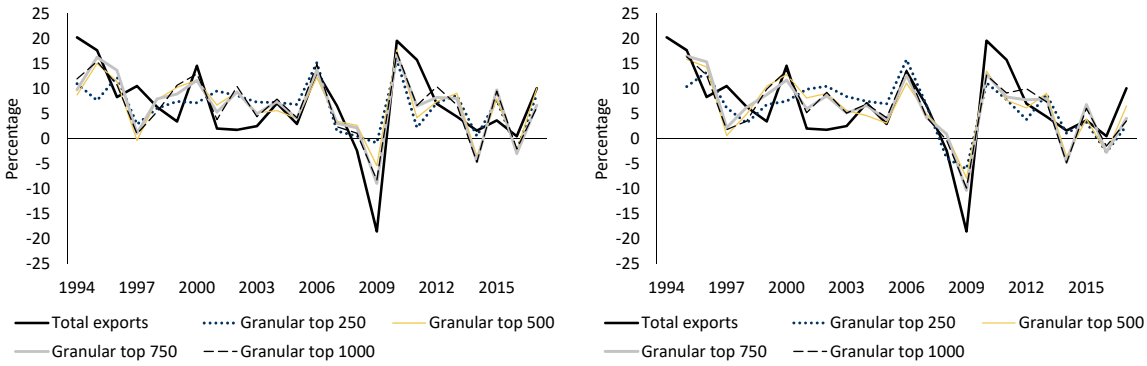

(A) Current values of the granular residual

(B) Current and lagged values of the granular residual

FIGURE 5: Actual and fitted values of the aggregate growth of Portuguese exports

Notes: The figure illustrates the estimation results of Equation (5). Panel a) plots the observerd rate of change of total Portuguese exports and the predicted values using the current value of the granular residuals of the top 250, 500, 750 and 1000 exporting firms. In Panel b), the fitted values are obtained using both the current and lagged values of the granular residual for the four groups of firms.

6. Concluding remarks

In recent years, there has been a significant improvement on the understanding of the micro-origins of aggregate fluctuations. When the firm-size distribution is very fat-tailed, idiosyncratic shocks to the largest firms directly contribute to aggregate dynamics. Hence, the role of the top units in a distribution links with the concept of granularity of Gabaix (2011).

This article shows that the tail of the distribution of Portuguese firm-level exports of goods is very heavy and adequately described by a power law with an exponent close to 1. Empirically, the existence of a power law distribution of firm size with such a low exponent indicates that Portuguese exports are very concentrated in a few large firms.

Moreover, the concentration of Portuguese export flows in the top firms tends to increase over time, in particular until 2009.

Using the empirical strategy of the granular residual of Gabaix (2011), we find that idiosyncratic shocks to the largest firms are relevant for total export growth over time, accounting for more than one-third of aggregate fluctuations. This means that volatility at the firm-level can affect aggregate export dynamics. Our findings for exports are in line with the results first discovered by Gabaix (2011): if firm-sizes in an economy are described by a Pareto distribution, then independent firm-level shocks can generate macroeconomic fluctuations, in accordance with the granular hypothesis.

This granular hypothesis has implications for monitoring and forecasting Portuguese exports. If a significant component of the dynamics of total exports originates from a small number of firms, it is essential to learn more about these top firms and the idiosyncratic shocks they are subject to. Besides studying macro shocks, monitoring the quantitative and qualitative information regarding a panel of large players might help in explaining and predicting aggregate export behaviour. These results have also policy implications, as small policy changes can have significant aggregate effects if they change the behaviour of top firms.

References

Amador, João and Luca David Opromolla (2013). "Product and destination mix in export markets." *Review of World Economics*, 149(1), 23–53.

Amador, João and Luca David Opromolla (2017). "Trade Margins and Cohorts of Traders in Portugal." *Banco de Portugal Economic Studies*, III(4), 1–18.

Axtell, Robert L. (2001). "Zipf Distribution of U.S. Firm Sizes." *Science*, 293(5536), 1818–1820.

Banco de Portugal (2016). "Portuguese international traders: some facts about age, prices and markets." *Economic Bulletin, Special Issue*, October 2016, 101–115.

Bastos, Paulo, Daniel A. Dias, and Olga A. Timoshenko (2018a). "Learning, prices and firm dynamics." *Canadian Journal of Economics*, 51(4), 1257–1311.

Bastos, Paulo and Joana Silva (2010). "The quality of a firm's exports: Where you export to matters." *Journal of International Economics*, 82(2), 99–111.

Bastos, Paulo and Joana Silva (2012). "Networks, firms, and trade." *Journal of International Economics*, 87(2), 352–364.

Bastos, Paulo, Joana Silva, and Eric Verhoogen (2018b). "Export Destinations and Input Prices." *American Economic Review*, 108(2), 353–392.

Behrens, Kristian, Gregory Corcos, and Giordano Mion (2013). "Trade Crisis? What Trade Crisis?" *The Review of Economics and Statistics*, 95(2), 702–709.

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott (2009). "Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods." In *Producer Dynamics: New Evidence from Micro Data*, edited by Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, NBER Chapters, chap. 14, pp. 513–552. National Bureau of Economic Research (NBER).

Békés, Gábor, László Halpern, Miklós Koren, and Balázs Muraközy (2011). "Still standing: how European firms weathered the crisis - The third EFIGE policy report." Blueprint Series 15, Bruegel.

Blanco-Arroyo, Omar and Simone Alfarano (2017). "Granularity of the Business Cycle Fluctuations: The Spanish Case." *Economía coyuntural, Revista de temas de coyuntura y perspectivas*, 2(1), 31–58.

Blanco-Arroyo, Omar, Alba Ruiz-Buforn, David Vidal-Tomás, and Simone Alfarano (2018). "On the determination of the granular size of the economy." *Economics Letters*, 173, 35–38.

Bricongne, Jean-Charles, Lionel Fontagné, Guillaume Gaulier, Daria Taglioni, and Vincent Vicard (2012). "Firms and the global crisis: French exports in the turmoil." *Journal of International Economics*, 87(1), 134–146.

Carvalho, Vasco M. and Basile Grassi (2019). "Large Firm Dynamics and the Business Cycle." *American Economic Review*, 109(4), 1375–1425.

Clauset, A., C. Shalizi, and M. Newman (2009). "Power-Law Distributions in Empirical Data." *SIAM Review*, 51(4), 661–703.

Crozet, Matthieu, Guy Lalanne, and Sandra Poncet (2013). "Wholesalers in international trade." *European Economic Review*, 58, 1–17.

da Silva, Sergio, Raul Matsushita, Ricardo Giglio, and Gunther Massena (2018). "Granularity of the top 1,000 Brazilian companies." *Physica A: Statistical Mechanics and its Applications*, 512(C), 68–73.

del Rosal, Ignacio (2013). "The granular hypothesis in EU country exports." *Economics Letters*, 120(3), 433–436.

del Rosal, Ignacio (2018). "Power laws in EU country exports." *Empirica*, 45(2), 311–337.

di Giovanni, Julian and Andrei A. Levchenko (2012). "Country Size, International Trade, and Aggregate Fluctuations in Granular Economies." *Journal of Political Economy*, 120(6), 1083–1132.

di Giovanni, Julian and Andrei A. Levchenko (2013). "Firm entry, trade, and welfare in Zipf's world." *Journal of International Economics*, 89(2), 283–296.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean (2014). "Firms, Destinations, and Aggregate Fluctuations." *Econometrica*, 82(4), 1303–1340.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean (2017). "Large Firms and International Business Cycle Comovement." *American Economic Review*, 107(5), 598–602.

di Giovanni, Julian, Andrei A. Levchenko, and Isabelle Mejean (2018). "The Micro Origins of International Business-Cycle Comovement." *American Economic Review*, 108(1), 82–108.

di Giovanni, Julian, Andrei A. Levchenko, and Romain Rancière (2011). "Power laws in firm size and openness to trade: Measurement and implications." *Journal of International Economics*, 85(1), 42–52.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2011). "An Anatomy of International Trade: Evidence From French Firms." *Econometrica*, 79(5), 1453–1498.

Ebeke, Christian H and Kodjovi M. Eklou (2017). "The Granular Origins of Macroeconomic Fluctuations in Europe." IMF Working Papers 17/229, International

Monetary Fund (IMF).

Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch (2012). *Modelling Extremal Events: for Insurance and Finance*. Springer.

Eppinger, Peter S., Nicole Meythaler, Marc-Manuel Sindlinger, and Marcel Smolka (2018). "The great trade collapse and the Spanish export miracle: Firm-level evidence from the crisis." *The World Economy*, 41(2), 457–493.

Esteves, Paulo Soares, Miguel Portela, and António Rua (2018). "Does domestic demand matter for firms' exports?" Working Paper 26-2018, Banco de Portugal.

Freund, Caroline and Martha Denisse Pierola (2015). "Export Superstars." *The Review of Economics and Statistics*, 97(5), 1023–1032.

Friberg, Richard and Mark Sanctuary (2016). "The contribution of firm-level shocks to aggregate fluctuations: The case of Sweden." *Economics Letters*, 147(C), 8–11.

Fujiwara, Yoshi, Corrado Di Guilmi, Hideaki Aoyama, Mauro Gallegati, and Wataru Souma (2004). "Do Pareto–Zipf and Gibrat laws hold true? An analysis with European firms." *Physica A: Statistical Mechanics and its Applications*, 335(1), 197–216.

Gabaix, Xavier (2009). "Power Laws in Economics and Finance." *Annual Review of Economics*, 1(1), 255–294.

Gabaix, Xavier (2011). "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79(3), 733–772.

Gabaix, Xavier (2016). "Power Laws in Economics: An Introduction." *Journal of Economic Perspectives*, 30(1), 185–206.

Gabaix, Xavier and Rustam Ibragimov (2011). "Rank - 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents." *Journal of Business & Economic Statistics*, 29(1), 24–39.

Gao, Baojun, Wai Kin (Victor) Chan, and Hongyi Li (2015). "On the increasing inequality in size distribution of China's listed companies." *China Economic Review*, 36(C), 25–41.

Gaubert, Cecile and Oleg Itskhoki (2018). "Granular Comparative Advantage." NBER Working Papers 24807, National Bureau of Economic Research (NBER).

Gnocato, Nicolò and Concetta Rondinelli (2018). "Granular sources of the Italian business cycle." Working paper 1190, Bank of Italy.

Hill, Bruce M. (1975). "A Simple General Approach to Inference About the Tail of a Distribution." *The Annals of Statistics*, 3(5), 1163–1174.

Karasik, Leonid, Danny Leung, and Ben Tomlin (2016). "Firm-Specific Shocks and Aggregate Fluctuations." Staff Working Papers 16-51, Bank of Canada.

Mayer, Thierry and Gianmarco Ottaviano (2008). "The Happy Few: The Internationalisation of European Firms." *Intereconomics: Review of European Economic Policy*, 43(3), 135–148.

Mion, Giordano and Luca David Opromolla (2014). "Managers' mobility, trade performance, and wages." *Journal of International Economics*, 94(1), 85–101.

Miranda-Pinto, Jorge and Yuanting Shen (2019). "A Granular View of the Australian Business Cycle." *Economic Record*, 95(311), 407–424.

Nagengast, A. J. (2019). "An N-dimensional generalization of the Amiti–Weinstein estimator." *Applied Economics Letters*, 26(8), 669–676.

Newman, M. E. J. (2005). "Power laws, Pareto distributions and Zipf's law." *Contemporary Physics*, 46(5), 323–351.

Nicolau, João and Paulo M. M. Rodrigues (2019). "A New Regression-Based Tail Index Estimator." *The Review of Economics and Statistics*, 101(4), 667–680.

Okuyama, K, M Takayasu, and H Takayasu (1999). "Zipf's law in income distribution of companies." *Physica A: Statistical Mechanics and its Applications*, 269(1), 125–131.

Wagner, Joachim (2012). "The German manufacturing sector is a granular economy." *Applied Economics Letters*, 19(17), 1663–1665.

Wagner, Joachim (2013). "The granular nature of the great export collapse in German manufacturing industries, 2008/2009." *Economics - The Open-Access, Open-Assessment E-Journal*, 7, 1–21.

Zhang, Jianhua, Qinghua Chen, and Yougui Wang (2009). "Zipf distribution in top Chinese firms and an economic explanation." *Physica A: Statistical Mechanics and its Applications*, 388(10), 2020–2024.

Non-technical summary

April 2020

The measurement of labour market slack: An empirical analysis for Portugal

Fernando Martins, Domingos Seward

This paper presents a perspective on the measurement of labour market slack in Portugal taking into account the degree of attachment of several groups of non-employment. We discuss the adequacy of the conventional criteria used for the measurement of unemployment. Following the relevant strand of literature on this topic, we apply a classification of labour market status based on the transition behaviour of non-employed individuals. In this context, we provide a comprehensive assessment of the heterogeneity in the Portuguese labour market

The results presented are based on the Portuguese Labour Force Survey. We observe that the unemployed have a higher probability of finding a job than the marginally-attached – those inactive individuals which express desire to work –, who, in turn, move to employment with a probability roughly 10 percentage points above that of the non-attached – those inactive individuals who do not express a desire to work (Table 1). On the basis of the statistical tests of equivalence, we conclude that the marginally-attached group constitutes a distinct state in the labour market. Moreover, marginally-attached workers display a transition behaviour closer to unemployment than to the non-attached.

We also find significant heterogeneity among the marginally-attached. In particular, the subgroup which reports "waiting" as a reason for not having searched a job displays a much higher transition rate into employment as well as a lower probability of moving to non-attachment. Even though the performed statistical tests for the equivalence of these groups lead to the rejection of their equivalence, one may argue that this is mainly driven by the fact that the "waiting" subgroup exhibits a much stronger attachment to the labour force than the unemployed. Within the marginally-attached population, we also observe that those individuals who have searched for a job but are still classified as inactive display a transition rate into employment which is comparable to the long-term unemployed, even after controlling for individual characteristics.

In addition, we find substantial heterogeneity among the non-attached. This is due to the fact that the so-called long-term future job starters display the highest degree of attachment to the labour market, judging by its high average transition rate into employment (36.4%). Although these individuals also frequently move to non-attachment (33.3%), their transition behaviour is closer to unemployment than to the rest of non-employed. Therefore, its classification as inactive might not be adequate.

Status	Stocks (thousands)	Flows into employment (%)	Transition rates into employment (%)
1. Unemployment	445.3	65.0	19.8
Short-term unemployed	203.0	38.1	25.7
Long-term unemployed	242.1	26.8	14.3
2. Marginally-attached (want)	208.4	20.0	13.4
Inactive searcher	14.0	1.7	16.4
Waiting	23.1	4.5	31.0
Discouraged	55.8	4.4	9.7
Personal reasons	47.1	3.3	9.3
Other reasons	68.3	6.1	13.2
3. Non-attached (do not want)	2,468.4	55.9	3.0
(Long-term) future job starts	1.4	0.5	36.4
Other non-attached	2,467.2	55.6	3.0

TABLE 1. Summary statistics by labour market groups, 1998:1-2019:3

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The values are the quarterly averages from 1998:1 to 2019:3. The observations from 2010 to 2011 are not considered in the calculations to avoid possible effects resulting from the methodological change of the LFS in 2011:1.

These results suggest possible shortcomings in those analyses which use slack measures based strictly on the job search criterion, as it is the case of the unemployment rate. A broader analysis of the labour market seems to be appropriate for an accurate assessment of the labour market slack.

The measurement of labour market slack: An empirical analysis for Portugal

Fernando Martins
Banco de Portugal
Universidade Lusíada de Lisboa
REM - Research in Economics and Mathematics, UECE

Domingos Seward
Banco de Portugal

April 2020

Abstract
This paper presents a perspective on the measurement of labour market slack in Portugal based on the degree of attachment of several groups of non-employed individuals. We discuss the adequacy of the conventional criteria used for the measurement of unemployment. Following the relevant strand of literature on this topic, we apply a classification of labour market status based on the transition behaviour of non-employed individuals. We conclude that some subgroups in inactivity display a transition behaviour into employment which is closer to the unemployed. This suggests that the classification of some individuals as inactive might not be adequate, since they show considerable attachment to the labour market and we reject the equivalence relative to their inactive counterparts. (JEL: C81, E24, J20)

1. Introduction

The recovery experienced by the Portuguese economy since 2013 has occurred in parallel with a substantial improvement in labour market conditions. The unemployment rate has decreased significantly, reaching the levels observed in 2003, while employment has increased back to pre-crisis levels. In spite of the improving labour market conditions, wage growth remains below the levels one would expect considering the cyclical position of the economy and the decrease in labour market slack. Indeed, whereas the unemployment rate has followed a decreasing

Acknowledgements: The authors wish to thank David Neves, António Antunes, João Amador, Lucena Vieira, Nuno Alves, Pedro Gomes, Pedro Portugal, Pedro Raposo, and the editor Pedro Duarte Neves for valuable inputs, comments, and suggestions. The authors also thank the comments by participants at the 7^{th} UECE Conference on Economic and Financial Adjustments, at the Lisbon School of Economics and Management (ISEG), Universidade de Lisboa, as well as the "Exchange" seminar, at the Economics and Research Department, Banco de Portugal. Part of this work was developed while Domingos Seward was a student at ISEG. The authors also thank Statistics Portugal for providing access to the microdata from the Portuguese labour force survey. The analyses, opinions, and findings expressed in this article are those of the authors and do not necessarily coincide with those of Banco de Portugal or the Eurosystem. UECE – Research Unit on Complexity and Economics is supported by Fundação para a Ciência e a Tecnologia (FCT).
Note: This article was submitted for approval on February 21, 2020, before the World Health Organization declared the COVID-19 a pandemic on March 11.

E-mail: fmartins@bportugal.pt; dseward@bportugal.pt

path for several years, wage growth remains lower than it was before the global financial crisis for equivalent levels of unemployment (OECD, 2018). This background of low unemployment and modest wage inflation has been referred to as an economic "puzzle". One the explanations pointed out is related with the inability of the headline unemployment rate to capture the true extent of labour market slack (see, *e.g.* Yellen, 2014).

The evidence shows that the non-employed population appears to be heterogeneous. While the distinction between the employed and the non-employed is straightforward, the boundary between the unemployed and the inactive is difficult to trace. For example, some persons classified as inactive can be considered close to unemployment if they have recently searched for a job or if they express desire to work. On the other hand, other inactive persons show little or no attachment to the labour force, namely by expressing no desire to work. Most of these individuals are less likely to find a job compared to those who have recently become unemployed, but the examination of longitudinal data on worker flows suggests that some subgroups within inactivity are at least as likely to find a job as the unemployed. Moreover, although the chance of transitioning from inactivity to employment is on average lower than it is from unemployment, the comparatively large size of the inactive population implies that these transitions can contribute substantially to the growth in employment, especially when unemployment decreases during periods of economic expansion. As discussed by Jones and Riddell (1999, 2006), one important implication for economic policy is that any effort towards measuring the degree of slack in the labour market by dichotomising the non-employed population into "unemployment" and "inactivity" is unable to comprehensively capture the complexity of labour market dynamics.

This paper presents a perspective on the measurement of labour market slack in Portugal based on the degree of attachment of several groups of individuals in the labour market. In this context, we propose an allocation of individuals across the three conventional states (employment, unemployment, and inactivity). We also discuss the adequacy of the conventional criteria used for the measurement of unemployment. Following the relevant strand of literature on this topic, we apply a classification of labour market status based upon the information on the transition behaviour of non-employed individuals.

2. The dataset

The results presented are based upon the Portuguese Labour Force Survey[1] (LFS), a household survey conducted quarterly by Statistics Portugal (*Instituto Nacional de Estatística*), with the goal of characterising the Portuguese labour market.

The LFS collects individual information on several features pertaining to the labour market, as well as demographic and socio-economic characteristics of the respondents. On the basis of this information, Statistics Portugal provides quarterly estimates for the stocks of employment, unemployment, and inactivity, which in turn are used for computing several indicators, such as the unemployment rate.

Each quarter, Statistics Portugal surveys approximately 22,000 households. The total sample is composed of six sub-samples which follow a rotation scheme, whereby each quarter 1/6 of the sample is rotated out and the remaining 5/6 are kept on the sample. Thus, once selected into the LFS sample, households should be interviewed for six consecutive quarters. Considering this feature of the sample, one can observe the labour force status for 5/6 of the respondents included in the sample in adjacent quarters, which enables the computation of worker flows and transition rates across states.[2]

In this article, we have used the LFS microdata for the period ranging from 1998:1 to 2019:3.

3. The measurement of the degree of labour market slack

The unemployment rate is the most commonly used measure of labour market slack. It is defined as the ratio between the number of unemployed individuals over the labour force. Labour Force Surveys (LFS) constitute the main source of data for the estimation of the number of unemployed individuals and their characterisation. Labour market statistics split the working-age population into three mutually-exclusive groups: the employed, the unemployed, and the inactive (*i.e.* the group of individuals deemed to be out of the labour force). However, while the distinction between the employed and the non-employed is straightforward, the boundary between the unemployed and the inactive is less clearcut.

According to the Portuguese LFS, which follows the general guidelines set by the International Labour Organization (ILO), an unemployed person must fulfil three criteria: (i) did not work during the reference week, (ii) is available to work during the reference week or within the next two weeks, and (iii) has actively searched for work during the reference week or within the previous three weeks.[3] The classification relies on the degree of attachment to the labour market, which is based crucially on the

1. In Portuguese, *Inquérito ao Emprego*.

2. See INE (2015) for additional information regarding the methodology of the LFS and the statistical inference procedure used for the construction of the sample. For a detailed description of the method used for calculating the gross worker flows and the transition rates, see Martins and Seward (2019, p. 34–35).

3. Those individuals which have not searched but are due to start work in the next three months are also included.

job search criterion. However, such a requirement may not be sufficient to completely capture the degree of slack contained in the labour market. The job search criterion is usually not defined with respect to time or pecuniary inputs and, importantly, it does not refer to the characteristics of the job, *e.g.* the offered wage (Jones and Riddell, 1999, 2006).

In addition, the non-employed population seems to be a particularly heterogeneous group. By adopting a classification of labour market status using the information obtained from the LFS (Table 1), we conclude that some persons classified as inactive can be considered close to unemployment if, for example, they have recently searched for a job, if they express desire to work, or if they are about to start a new job but beyond the three-month threshold for an individual to be classified as unemployed. On the other hand, other inactive individuals show little or no attachment to the labour market, either because they display little marketable skills or because they do not desire to work. A group classified as inactive which has been the subject of increasing policy analysis is the so-called marginally-attached individuals, which are those individuals that want to work, but who are not actively searching for a job. Within the group of marginally-attached individuals, one can single out the group of discouraged workers, which comprises those individuals which want to work, but do not report actively searching for a job due to economic reasons (Table 1).[4]

In this context, unemployment may not be a sufficient metric to assess the degree of labour market slack if the requirements do not sort individuals appropriately relative to their willingness to work and/or their chance of finding a job, *i.e.* if considerable fractions of the non-employed population which do not satisfy the requirements to be classified as unemployed, would answer in a similar way when finding a given job vacancy.[5]

In practice, many non-employed persons become employed without being recorded in unemployment. Table 2 summarises the average quarterly employment inflows disaggregated by several subgroups of origin, over the period from 1998:1 to 2019:3.[6] We observe that indeed employment inflows originating from inactivity are substantial, and represent on average 76 thousand individuals each quarter. This figure compares with an average of 65 thousand individuals originating from unemployment. In particular, we observe that the transition pattern differs considerably between the inactivity subgroups. For example, 13.4% of the marginally-attached (those that express the desire to work) move into employment each quarter on average. On the other hand, the non-attached workers (those that do not desire to work) are much less likely to move into employment (3.0%). Still, given the considerable size of the non-attached, such a low transition rate translates into non-negligible gross flows into employment in absolute terms (55.9 thousand each quarter on average). In addition, differences among the

4. Among the reasons for not actively searching for work is the belief that no work is available.

5. See, *e.g.* Jones and Riddell (1999, 2006) and Schweitzer (2003) for interesting discussions on this issue.

6. The LFS is based upon a probabilistic sample. Therefore, small aggregates tend to be estimated with lower precision. In this context, the values presented in Table 2 should be interpreted with some caution.

unemployed are noteworthy. As expected, the short-term unemployed are much more likely to move into employment than the long-term unemployed (25.7% *versus* 14.3%).

Such apparent differences in the employability of the above-mentioned labour market groups are also associated to differences in some socio-demographic characteristics (Table 3). In this context, we observe that the marginally-attached exhibit levels of education considerably higher than the non-attached and much closer to the unemployed (only 4.8% of the marginally-attached have no education, which compares with a value of 16.7% for the non-attached, and of 3.3% for the unemployed). Likewise, 23.0% of the marginally-attached are included in the 25 to 34 years-old cohort *vis-à-vis* 4.7% of the non-attached. We also observe that the marginally-attached tend to be on average less time out of employment in comparison with the non-attached (on average 15.4% of the marginally-attached are out of employment for less than a year, whereas only 2.3% of the non-attached display periods out of employment of less than a year).

Regarding the wages earned once these individuals become employed, the marginally-attached report a median net wage comparable to that of the unemployed and very close to the one reported by the long-term unemployed.[7] On the other hand, the reported median net wage by those inactive individuals that do not express desire to work is considerably lower, which is a further indication of heterogeneity among the inactive population.

We also notice that three inactive subgroups are particularly relevant in terms of their estimated transition rates into employment: the inactives who search for work, the inactives who report “waiting" as a reason for not having searched, and the inactives who are due to start a job in more than three months. These subgroups also display sociodemographic characteristics which are closer to the unemployed, distinguishing them from the remaining inactives. In particular, these subgroups exhibit higher levels of education in comparison to the remaining inactives, they are younger (with particular relevance of the 25 to 34 years-old cohort), they remain out of employment for less time, and they find on average a higher proportion of open-ended contracts, as well as higher entry net wages.

The aforementioned groups of inactivity are quantitatively relevant and could thus affect the perspective on the amount of underutilised labour supply in the market. In the data, the marginally-attached represent, on average, 6.5% of the non-employed population, whereas the discouraged workers account for 1.7% (Table 2). Whereas most individuals in these groups have a lower chance of moving to employment compared with the recently unemployed, they often obtain work. Therefore, they may serve to enlarge the pool of unemployed as an important potential source of labour supply.

7. See Martins and Seward (2019, p. 37) for a detailed analysis of the reported median net wages once the non-employment individuals find a job.

Status	Abbreviation	Adopted definition
1. Unemployment	U	Non-employed individuals who did not work during the reference week, are available to work during the reference week or within the next two weeks, and have actively searched for work during the reference week or within the previous three weeks (the individuals which have not searched but are due to start work in the next three months are also included).
Short-term unemployment	$U(ST)$	Unemployed for less than 12 months.
Long-term unemployment	$U(LT)$	Unemployed for more than 12 months.
2. Marginally-attached	M	Inactive individuals who to want work.
Inactives who search	$M(S)$	Inactive individuals who search for work (includes individuals which search passively).
Inactives who wait	$M(W)$	Inactive individuals who report “waiting" as a reason for not searching for work (includes workers in temporary layoff).
Discouraged	$M(D)$	Inactives who report economic reasons not searching for work (believe there are no jobs available, are too young/old, do not have enough education, do not know how to find a job or consider that it is not worth searching).
Personal reasons	$M(P)$	Inactives who report personal reasons for not searching for work (ilness or inability, need to take care after children/disabled/elderly or other personal reasons).
Other marginally-attached	$M(O)$	Inactives who report studying or training, retirement, and other reasons for not searching for work.
3. Non-attached	N	Inactives who do not want (nor search for) work.
Long-term future job starters	$N(FJS)$	Inactives who will start a job in more than three months (or within three months but not available to start work during the reference week or the following two weeks).
Other non-attached	$N(O)$	Other non-attached individuals (includes students, retired workers, domestic workers, disabled individuals and others).

TABLE 1. Definitions of the selected non-employed subgroups

Status	Stocks (thousands)	Fraction of non-employment (%)	Flows into employment (thousands)	Fraction of new employment created (%)	Transition rates into employment (%)
1. Unemployment	445.3	14.0	65.0	46.7	19.8
Short-term unemployed	203.0	6.4	38.1	28.8	25.7
Long-term unemployed	242.1	7.5	26.8	17.9	14.3
2. Marginally-attached (want)	208.4	6.5	20.0	13.4	13.4
Inactive searcher	14.0	0.4	1.7	1.1	16.4
Waiting	23.1	0.8	4.5	4.0	31.0
Discouraged	55.8	1.7	4.4	2.8	9.7
Personal reasons	47.1	1.5	3.3	2.1	9.3
Other reasons	68.3	2.1	6.1	3.6	13.2
3. Non-attached (do not want)	2,468.4	79.5	55.9	39.9	3.0
(Long-term) future job starters	1.4	0.1	0.5	0.4	36.4
Other non-attached	2,467.2	79.5	55.6	39.6	3.0

TABLE 2. Summary statistics for selected non-employed subgroups, 1998:1-2019:3

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The values are the quarterly averages from 1998:1 to 2019:3. The observations from 2010 to 2011 are not considered in the calculations to avoid possible effects resulting from the methodological change of the LFS in 2011:1.

	Unemployed			Marginally-attached			Non-attached	
	Total	Short-term Unemployed	Long-term Unemployed	Total	*of which:* Search for work	Waiting	Total	*of which:* Future job starters[(1)]
by Education								
No education	3.3	2.8	3.7	4.8	4.2	3.3	16.7	6.1
Basic education	63.4	59.1	67.1	68.1	61.9	64.2	64.8	55.8
Secondary education	20.2	22.3	18.5	18.0	22.0	18.2	13.6	23.5
Higher education	13.1	15.7	10.7	9.0	11.9	14.3	4.9	14.6
by Age Cohort								
15-24	22.3	30.8	14.1	25.7	23.8	21.7	27.7	30.2
25-34	26.9	30.0	24.1	23.0	28.3	27.9	4.7	26.9
35-44	21.4	19.5	23.4	18.6	20.3	22.4	5.0	16.7
$\geq$45	29.4	19.7	38.4	32.8	27.6	28.0	62.7	26.2
by Gender								
Female	52.9	53.0	53.2	62.4	61.5	56.9	60.2	48.3
Male	47.1	47.0	46.8	37.6	38.5	43.1	39.8	51.7
by Time Out of Employment								
<1 year	22.3	–	–	15.4	20.0	30.7	2.3	24.7
1-2 years	27.3	–	–	19.9	25.7	27.0	4.7	25.7
$\geq$2 years	50.4	–	–	64.7	54.4	42.3	93.0	49.6
by Type of Contract[(2)]								
Open-ended contract	13.6	12.7	15.3	16.2	17.9	17.7	15.9	26.3
Other types of contracts	86.4	87.3	84.7	83.8	82.1	82.3	84.1	73.7
Net wages once employed[(3)]	520.0	550.0	500.0	500.0	505.0	525.0	485.0	600.0

TABLE 3. Sociodemographic characterisation and reported median net wages by labour market groups, 1998:1-2019:3

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The values are the quarterly averages of the proportions (%) in each labour market group by sociodemographic characteristics in the period from 1998:1 to 2019:3. The observations from 2010 to 2011 are not considered in the calculations to avoid possible effects resulting from the methodological change of the LFS in 2011:1.
(1) Subgroup comprised by the inactive individuals which are due to start a job in more than three months.
(2) The values are the quarterly averages of the proportions (%) in each labour market group which either obtain work with an open-ended contract or with another type of contract (*e.g.* fixed-term contract), in the period from 1998:1 to 2019:3.
(3) The values are the reported median net wages in euros by each labour market group once the individuals transition into employment, in the period from 2012:1 to 2019:3.

4. Literature review

The study of the heterogeneity between labour market states is crucial for a comprehensive characterisation of the degree of labour market slack. The literature about transition rates between labour market states with implications for the classification of individuals was pioneered by Clark and Summers (1978). In analysing the dynamics of youth unemployment for the United States of America (USA), the authors claim that most of youth non-employment is not captured by the conventional unemployment statistics, since many stop searching and withdraw from the labour force. The distinction between unemployment and inactivity for youth might be meaningless if we consider the wide array of non-market options accessible to youths and the limitations imposed by unemployment compensation schemes on the eligibility of this group. The analysis suggests that the empirical distinction for this group between the above-mentioned statuses is considerably arbitrary and of little practical value.

More generally, Clark and Summers (1979) find that transitions between unemployment and employment in the USA are considerably lower in magnitude compared with transitions into and out of inactivity. In addition, many individuals appear to experience several changes in classification within a single non-employment spell, with repeated spells of unemployment discontinued by withdrawal from the labour force. Such evidence is supportive of a weak distinction between the unemployment and the inactive categories.

The findings obtained by these authors inspired several statistical analyses of the equivalence of the unemployment and inactivity categories. In their seminal article, Flinn and Heckman (1983) rationalised the distinction between labour market states based on transition probabilities. In this sense, individuals are said to belong to the same labour market state if they exhibit equivalent behaviour with respect to subsequent labour market status.[8] The authors proposed a statistical framework for testing the equivalence of labour market states in longitudinal data, based on a duration of status econometric approach. The authors test for the equivalence between the unemployment and inactive states for young white American males and find evidence that rejects this hypothesis.

Tano (1991) tests the hypothesis that unemployment and inactivity are behaviourally meaningless classifications using the Current Population Survey (CPS) gross flows data. To do so, the author employs a binary logit econometric framework. The results indicate that the two states are distinct for youth, whereas for prime-age individuals they are meaningless. In the same vein, Gönül (1992) extends the former analysis to a wider group of male and female highschool graduates, by employing a duration econometric model, with mixed results by gender.

Jones and Riddell (1999, 2006) extended the former literature by examining the transition behaviour within the unemployed and the inactive groups for the USA and Canada. The authors examine the equivalence between groups by applying multinomial

8. Therefore, two groups may be considered equally attached to the labour market if they are equally likely to move to employment in the following period.

and binary logit models for the transition behaviour of individuals. The authors find that the group of marginally-attached workers (comprising those inactives who do not search, but want work) constitutes a distinct state in the labour market. The authors discuss the adequacy of the unemployment measurement criteria, since they find that some subgroups in inactivity display a transition behaviour closer to unemployment.[9] Schweitzer (2003) obtains similar results for the United Kingdom.

Brandolini *et al.* (2006) also find evidence of substantial heterogeneity among the inactive group for European countries. The authors investigate the role of the four-week job search requirement by examining the behaviour of those individuals who search for work but did so more than four weeks before the survey interview. Their analysis is conducted by a non-parametric equality test. The results show that for most countries this group forms a distinct state in the labour market. In addition, the authors find that these individuals are equivalent to the unemployed when their last search effort was done not long before the four-week requirement, which highlights the arbitrariness of the criterion.

Centeno and Fernandes (2004) study the heterogeneity in the Portuguese labour market. The data used in their work is drawn from the Portuguese LFS for the period ranging from 1992:1 to 2003:4. The authors adopt a duration econometric framework to model transition probabilities. The results suggest that the marginally-attached group is a distinct labour market state in Portugal. These findings have been confirmed by Centeno *et al.* (2010), with implications for the Non-Accelerating Inflation Rate of Unemployment (NAIRU).

Regarding the heterogeneity within the unemployed state, Hornstein (2012) and Krueger *et al.* (2014) show that even within unemployment the behaviour of the long-term unemployed points towards considerable variations in employability. Likewise, other investigations conclude that the short-term and the long-term unemployed exhibit substantial differences in their transition behaviour to employment (see Kroft *et al.*, 2012 and Eriksson and Rooth, 2014).

5. Heterogeneity in the Portuguese labour market

5.1. Statistical framework

The adopted statistical framework folows the seminal contribution by Flinn and Heckman (1983), subsequently extended by Jones and Riddell (1999, 2006), by focusing on transition rates to assess the equivalence between states and the extent of heterogeneity in the labour market.

Let Y_t be a random variable describing the status of persons in the labour market at quarter t.[10] We assume that the transition of workers across labour market states

9. For instance, the subgroup comprising those individuals which report “waiting" as a reason for not having searched for a job.

10. For the purpose of this work, Y_t is assumed discrete and takes on values corresponding to k mutually-exclusive and exhaustive states.

is represented by a discrete Markov chain of order 1. Therefore, the data generating process, $\{Y_t\}_{t=1}^{T}$, follows:

$$\Pr(Y_t = i|Y_{t-1}, Y_{t-2}, \ldots, Y_1) = \Pr(Y_t = i|Y_{t-1})\,, \tag{1}$$

wherein $i = 1, 2, \ldots, k$ indexes the observed status in the Y_t domain. The process represented by equation (1) respects the Markov property.[11]

The probabilities of transition from state i to state j over quarters $t-1$ and t are given by:

$$p_{ij,t} = \Pr(Y_t = j|Y_{t-1} = i),\ i, j = 1, 2, \ldots, k\,. \tag{2}$$

We start by considering four labour market states ($k = 4$). The transitions across employment (E), unemployment (U), the group of the marginally-attached (M), and the group of the non-attached (N) is summarised by the four-by-four transition matrix P, where the ij^{th} element, p_{ij}, represents the probability of a person moving from state $i \in \{E, U, M, N\}$ in the current quarter to state $j \in \{E, U, M, N\}$ in the following quarter:

$$P_t = \begin{pmatrix} p_{EE} & p_{EU} & p_{EM} & p_{EN} \\ p_{UE} & p_{UU} & p_{UM} & p_{UN} \\ p_{ME} & p_{MU} & p_{MM} & p_{MN} \\ p_{NE} & p_{NU} & p_{NM} & p_{NN} \end{pmatrix}_t . \tag{3}$$

In this paper, we apply an evidence-based categorisation of labour market status by exploiting the information on the results of the transition behaviour of non-employed individuals.[12] Therefore, we classify individuals into the same state if they exhibit equivalent behaviour regarding subsequent state.[13]

5.2. Descriptive analysis

Table 4 shows the estimated quarterly transition rates for adjacent quarters averaged across the sample period. For transitions into E, there is a noticeable difference between U and M as origin states, with the transition rate from U at 19.8%, almost 6 percentage points above that of M (13.4%). Moreover, there is a striking difference between the M and N as origin groups, with the transition rate from N to E averaging only 3.0%. In addition, for each non-employment destination state, the transition rates between origin groups U and M and between M and N differ considerably: $\hat{p}_{MN} = 27.4\% > \hat{p}_{UN} = 8.3\%$ and $\hat{p}_{MU} = 22.8\% > \hat{p}_{NU} = 1.5\%$.

11. Thus, the observed values for Y_t depend only on the current status. Such assumption is rather strict and whenever possible should be tested. We have conducted a robustness analysis to this assumption by restricting the sample to those individuals in non-employment for less than twelve months and conclude that the findings hold.

12. The pioneering authors of this approach include Flinn and Heckman (1983) and Jones and Riddell (1999, 2006). See Centeno and Fernandes (2004) and Centeno *et al.* (2010) for applications to Portugal.

13. For instance, one may consider two groups to be equally attached to the labour market if they are equally likely to move to employment in the next period. The approach we take generalises this idea to all the statuses considered.

From the examination of the transition rates over the sample period (Figure 1), we conclude that (i) these exhibit, in general, considerable stability over time, with the exception of those into employment for which the cyclical pattern is very marked, (ii) the ordering of the transition rates is the same in every quarter over the sample period, with $\hat{p}_{UE} > \hat{p}_{ME} > \hat{p}_{NE}$, $\hat{p}_{UU} > \hat{p}_{MU} > \hat{p}_{NU}$, and $\hat{p}_{NN} > \hat{p}_{MN} > \hat{p}_{UN}$, and ($iii$) the difference between $\hat{p}_{UE}$ and $\hat{p}_{ME}$ is consistently much lower than the difference between $\hat{p}_{ME}$ and $\hat{p}_{NE}$. The fact that $\hat{p}_{ME}$ is close to $\hat{p}_{UE}$ is an indication that an expressed desire to work among non-employed individuals conveys substantial information about their attachment to the labour market.

In order to examine the extent of heterogeneity within the labour market states considered, we compute the average transition rates by detailed origin state (Table 4). We perform the conventional split of U by duration. The short-term unemployed are almost twice as likely to move into E (25.7%) relative to the long-term unemployed (14.3%). Conversely, the long-term unemployed have a higher chance to remain unemployed or to move into inactivity in the following quarter.

Furthermore, we find important heterogeneity within the M group. The striking result is that the "waiting" subgroup shows a transition rate into E (31.0%) considerably higher than the other subgroups, as well as a lower transition rate into N (13.9%). Moreover, the subgroup comprising those M individuals which report having searched for a job also displays significant attachment to the labour market, since 16.4% of these individuals move to employment each quarter on average, which is above the transition rate estimated for the long-term unemployed (14.3%).[14]

A measurement issue relates to those individuals who are not searching for work but who have found a job due to start more than three months after the survey interview[15]. We refer to these “officially" inactive but highly attached non-employed as long-term future job starters. In Portugal, as in many other countries, these non-employed are classified as inactive even though they display the largest estimated transition rate into E of all the subgroups considered in this study (36.4%).[16] In Portugal, this subgroup of N amounts to 1.4 thousand individuals each quarter on average (Table 2).

14. On the other hand, these non-employed also move to non-attachment quite often (20.0%); still, this figure is considerably below the corresponding transition rate displayed by most of their marginally-attached counterparts.

15. Those that are due to start a job within three months but who do not meet the availability criterion are also included.

16. However, they also exhibit a high transition rate into inactivity (33.3%), which makes it hard to evaluate this classification practice based on these data.

To / From	E	U	M	N
E	0.963	0.015	0.005	0.016
	(0.002)	(0.001)	(0.000)	(0.001)
U	0.198	0.625	0.094	0.083
	(0.005)	(0.007)	(0.003)	(0.004)
subgroups of U				
$U(ST)$	0.257	0.583	0.082	0.078
	(0.006)	(0.008)	(0.002)	(0.003)
$U(LT)$	0.143	0.665	0.103	0.089
	(0.004)	(0.006)	(0.004)	(0.004)
M	0.134	0.228	0.364	0.274
	(0.003)	(0.005)	(0.004)	(0.005)
subgroups of M				
$M(S)$	0.164	0.348	0.295	0.200
	(0.009)	(0.011)	(0.010)	(0.010)
$M(W)$	0.310	0.295	0.263	0.139
	(0.016)	(0.012)	(0.011)	(0.007)
$M(D)$	0.097	0.208	0.436	0.259
	(0.004)	(0.005)	(0.006)	(0.006)
$M(P)$	0.093	0.182	0.393	0.331
	(0.005)	(0.007)	(0.007)	(0.008)
$M(O)$	0.132	0.213	0.325	0.330
	(0.006)	(0.008)	(0.009)	(0.010)
N	0.030	0.015	0.029	0.926
	(0.002)	(0.000)	(0.003)	(0.004)
subgroups of N				
$N(FJS)$	0.364	0.179	0.127	0.333
	(0.034)	(0.025)	(0.024)	(0.033)
$N(O)$	0.030	0.015	0.029	0.926
	(0.002)	(0.000)	(0.003)	(0.004)

TABLE 4. Average quarterly transition rates by labour market groups, 1998:1-2019:3

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The values are the quarterly averages from 1998:1 to 2019:3. The standard errors are in parentheses. The observations from 2010 to 2011 are not considered in the calculations to avoid possible effects resulting from the methodological change of the LFS in 2011:1.

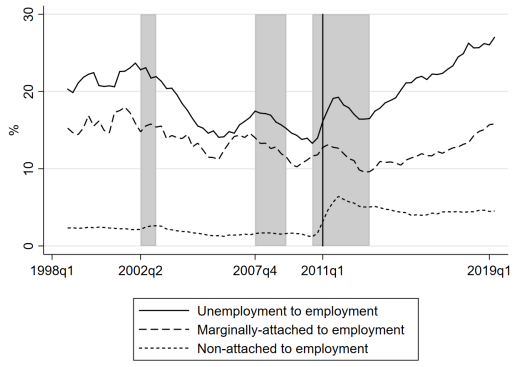

(A) Transition rates into employment

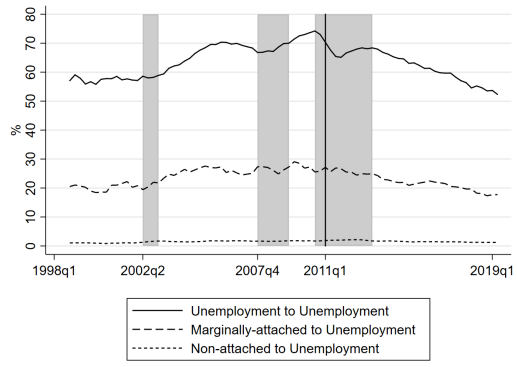

(B) Transition rates into unemployment

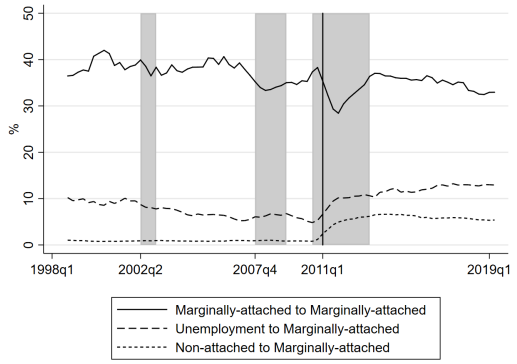

(C) Transition rates into marginal-attachment

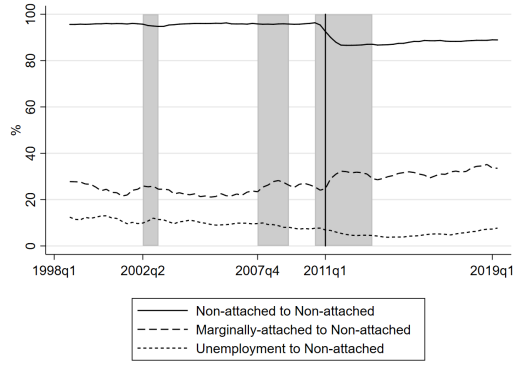

(D) Transition rates into non-attachment

FIGURE 1: Average quarterly transition rates by groups of inactivity, 1998:1-2019:3

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The series are a four-quarter moving-average to abstract from the marked seasonal pattern. The shadings indicate recessions according to Rua (2017). The vertical line signals the methodological change of the LFS in 2011:1.

5.3. Conditional assessment

5.3.1. Econometric model

The average transition rates analysed so far consider individuals which differ on various characteristics. Therefore, it is crucial to assess whether the findings are essentially driven by compositional effects.[17] For this purpose, we estimate multinomial logit models (MLM) on the determinants of transitions across several labour market states.[18][19]

In logistic models, the probability of moving into a category is compared to the probability of being in the baseline category. Considering that, in the most general case, we have k categories, such an approach requires the computation of $k-1$ equations, one for each destination state with respect to the baseline. Thus, there will be $k-1$ predicted log-odds. If we define the j^* as the baseline outcome, we obtain the following system:[20]

$$f_j(\boldsymbol{x}_{h,t}) = \ln \frac{\Pr(Y_{h,t}=j|Y_{h,t-1}=i,\boldsymbol{x}_{h,t})}{\Pr(Y_{h,t}=j^*|Y_{h,t-1}=i,\boldsymbol{x}_{h,t})} = \alpha_j + \boldsymbol{x}'_h\boldsymbol{\beta}_j,\ j \neq j^*, \tag{4}$$

where h indexes the person, $Y_{h,t}$ denotes the first-order Markov chain for person h in period t, $\boldsymbol{x}_{h,t}$ refers to the vector of conditioning individual characteristics, α_j denotes a constant, and $\boldsymbol{\beta}_j$ denotes the vector of regression coefficients.

We aim at testing for the equivalence between the probabilities of transitions into different labour market states, for instance, to test whether one can pool individuals originating from state M with individuals originating from state N. To do so, we take all individuals in the sample who belong to states M or N state in the first period, such that their three possible destinations are E, U or to remain in inactivity ($M+N$), and we estimate a multinomial logit regression (equation (4)).[21] The covariates refer to the personal and socio-economic characteristics of the respondent, as well as to a set of seasonal and regional dummy variables.[22] Afterwards, an unrestricted model is estimated by adding a dummy variable identifying the individuals who were originally in state M. The origin state dummy is interacted with each explanatory variable. In order to test for the equivalence between M and N, we employ a likelihood-ratio test, under

17. Such that different types of persons are more or less likely to belong to different groups than others (with an impact on the respective transition rates) or whether the findings still hold after controlling for such differences.

18. As opposed to Jones and Riddell (1999, 2006), we report the results from pooled multinomial logit regressions, since it has the advantage of increasing the sample size.

19. Such models will enable us to test whether two origin states are equivalent after controlling for the observable characteristics of the individuals.

20. In our model, we set $j^* = i$, *i.e.* we define the baseline outcome as the individual remaining in the previous state.

21. The models are estimated via the maximum likelihood procedure.

22. See Martins and Seward (2019, p. 41) for a detailed description of the covariates used in the regressions.

the null hypothesis that the groups are equivalent. The same reasoning is applied to test for the equivalence between other non-employed groups.[23]

5.3.2. Discussion of results

The results of the likelihood-ratio tests are reported in Table 5.[24]

	Time period	
H_0	1998:1-2010:4	2011:1-2019:3
$M = N$	9,021.85 (0.000)	5,720.34 (0.000)
$M = U$	2,816.79 (0.000)	6,759.55 (0.000)
$U = N$	121,896.63 (0.000)	84,699.43 (0.000)

TABLE 5. Likelihood-ratio tests in multinomial logit models for the equivalence between non-employment states

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The reported values are the observed likelihood-ratio test statistics for the respective H_0. The p-values are reported in parentheses.

We reject the statistical equivalence between M and N, M and U, and N and U, in all the periods.[25] Such results provide evidence supporting the hypothesis that the group comprising the marginally-attached individuals is distinct from the non-attached, as well as rejecting the equivalence between the marginally-attached and the unemployed. Furthermore, we reject the pooling of the unemployed and the non-attached groups. Hence, these formal statistical tests generally corroborate the evidence found for the empirical transition rates.

We also test the heterogeneity within each of the analysed groups (Table 6) and conclude that: (i) within the unemployed, our results point towards a rejection of the equivalence between the short-term and the long-term unemployed, (ii) within the marginally-attached, the statistical evidence leads to the rejection of the equivalence between its subgroups, and (iii) within the non-attached group, we test and reject the null of equivalence between the future job starters and the other non-attached.

Lastly, we conduct statistical tests for the equivalence between subgroups across the conventional classification criteria (Table 7). The tests again lead us to reject the equivalence between all the states considered. However, one can argue that, to the extent that the tests reject pooling the states, this is mainly due to the fact that, for example, the probability of the subgroup of inactive individuals which report "waiting" ($M(W)$)

23. Since the LFS was subject to a survey redesign in 2011:1 (INE, 2015; Neves, 2014), we conduct the tests separately for each survey.

24. Considering the interest in the equivalence tests rather than on the interpretation of the estimated regressions, for the sake of space we do not report the estimated regressions in this paper. See Martins and Seward (2019, p. 46–49) for the detailed results of the regressions.

25. This finding can be inferred from the large values for the observed likelihood-ratio test statistic and respective p-values which are equal to 0.000 for all the conducted tests.

	Time period	
H_0	1998:1-2010:4	2011:1-2019:3
$U(ST) = U(LT)$	3,415.48 (0.000)	2,508.15 (0.000)
$M(W) = M(S) = M(P) = M(D) = M(O)$	1,493.90 (0.000)	3,116.98 (0.000)
$N(FJS) = N(O)$	257.63 (0.000)	111.20 (0.000)

TABLE 6. Likelihood-ratio tests in multinomial logit models for the equivalence between subgroups of the same non-employment status

Source: Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The reported values are the observed likelihood-ratio test statistics for the respective H_0. The p-values are reported in parentheses. $U(ST)$, $U(LT)$, $M(S)$, $M(W)$, $M(D)$, $M(P)$, $M(O)$, $N(FJS)$, and $N(O)$ stand for short-term unemployed, long-term unemployed, marginally-attached searching, waiting, discouraged, personal reasons, other reasons, long-term future job starters, and other non-attached, respectively.

moving into employment is higher than that of the unemployed. The same conclusion can be inferred for those future job starters classified as inactive ($N(FJS)$).

	Time period	
H_0	1998:1-2010:4	2011:1-2019:3
$M(W) = U$	344.33 (0.000)	124.04 (0.000)
$M(S) = U$	137.56 (0.000)	389.24 (0.000)
$N(FJS) = U$	63.08 (0.002)	56.99 (0.004)

TABLE 7. Likelihood-ratio tests in multinomial logit models for the equivalence between subgroups of different non-employment statuses

Source: Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes: The reported values are the observed likelihood-ratio test statistics for the respective H_0. The p-values are reported in parentheses. U, $M(W)$, $M(S)$, and $N(FJS)$ stand for the unemployed, the marginally-attached searching, the marginally-attached waiting, and the long-term future job starters.

5.3.3. Limitations and robustness check

The MLM specification imposes the independence of irrelevant alternatives assumption (IIA) (Luce, 1959). Under this strong assumption, the relative probabilities of transitions into, *e.g.* E and U, would not change given the removal of the (irrelevant) alternative of transitions into inactivity. Such a scenario seems unrealistic considering that U is in several aspects closer to inactivity than to E.[26]

Hausman and McFadden (1984) developed a test of the IIA. We conduct this test for the multinomial logit models applied to test for the equivalence between M and N, M and U, and N and U. We obtain mixed results for each outcome depending on the category omitted from the full model.[27] Therefore, we cannot rule out the presence of the IIA in the multinomial logit models.

To assess the robustness of our results, we also estimate binary logit models, which can be viewed as imposing the polar assumption of complete dependence[28] (Table 8). A simpler model only controls for the seasonal and the regional pattern of the transitions rates, whilst a more complete model adds the usual demographic and socio-economic individual explanatory variables. By employing the equivalence tests to the simplest model, we find that three of the eleven inactive categories considered are in some sense comparable to unemployment with respect to their attachment to the labour market: the inactives who have searched for a job, the marginally-attached workers who are "waiting", and the long-term future job starters. The consideration of individual controls in a more complete model reinforces these findings, since the tests do not change significantly.

26. This modelling issue had been initially raised by Jones and Riddell (1999).

27. See the tables in Martins and Seward (2019, p. 50) for the detailed results of the Hausman and McFadden (1984) tests.

28. See Schweitzer (2003) for a similar application.

	Model with seasonal and regional controls	Model with individual controls
Subgroups of U		
U, short-term	6.415***	4.749***
	(0.101)	(0.103)
U, long-term	3.708***	2.693***
	(0.053)	(0.052)
Subgroups of M		
M, searched	3.527***▲▲	2.758***▲▲
	(0.156)	(0.128)
M, waiting	9.512***★★	7.023***★★
	(0.573)	(0.437)
M, discouraged	2.502***	2.030***
	(0.059)	(0.053)
M, personal reasons	2.221***	1.831***
	(0.061)	(0.055)
M, other reasons	2.306***	1.859***
	(0.049)	(0.050)
Subgroups of N		
N, future job starter	9.103***★★	6.591***★★
	(1.664)	(1.218)
N, student	0.753***	0.658***
	(0.013)	(0.018)
N, retired	*Excluded*	*Excluded*
N, domestic	1.343***	1.409***
	(0.023)	(0.028)
N, disabled	0.532***	0.453***
	(0.018)	(0.016)
N, other	2.213***	1.756***
	(0.055)	(0.047)
Number of observations	404,590	404,590
Pseudo-R^2	0.081	0.091

TABLE 8. Results of binary logit estimation and equivalence tests between non-employment subgroups

Source: Authors' calculations based on the Labour Force Survey (Statistics Portugal).
Notes:

1. The estimations are conducted on pooled data from the LFS for individuals occupying the non-employment group in the origin state;
2. The dependent variable takes on values 0 and 1, corresponding to two destination states: respectively, the pooled non-employment state and the employment state. Baseline category is remaining in the pooled non-employment state;
3. The reported values are the odd-ratios. The robust standard-errors are in parentheses;
4. Each model includes a constant. Seasonal and regional patterns are captured by dummy variables. The individual controls are: age, age squared, gender (male, female), marital status (single, married), education level (none, basic, secondary, higher);
5. The significance levels are: 10% (*); 5% (**); 1% (***);
6. The equivalence tests (at a 5% significance level) are: (★★) denotes coefficient statistically significantly greater relative to short-unemployed; (★) coefficient statistically significantly equal relative to short-unemployed; (▲▲) coefficient statistically significantly greater relative to long-unemployed; (▲) coefficient statistically significantly equal relative to long-unemployed.

6. Conclusion

This paper presents a perspective on the measurement of the level of slack in the Portuguese labour market, taking into acount the degree of attachment of several groups

in the labour market. In this context, we provide a comprehensive assessment of the heterogeneity in the Portuguese labour market, extending the work initially developed by Centeno and Fernandes (2004) and Centeno *et al.* (2010) to the non-employed population as a whole.

By disaggregating the non-employed population into three groups (the unemployed, the marginally-attached, and the non-attached), we find an evident distinction between each of these subgroups: the unemployed are more likely to move into employment than the marginally-attached, who, in turn, move to employment with a probability roughly 10 percentage points above that of the non-attached.[29] On the basis of the statistical tests of equivalence, we conclude that the marginally-attached group constitutes a distinct state in the labour market. Moreover, marginally-attached workers display a transition behaviour closer to the unemployed than to the non-attached.

We also find significant heterogeneity among the marginally-attached. In particular, the subgroup which reports “waiting" as a reason for not having searched displays a much higher transition rate into employment, as well as a lower probability of moving to non-attachment. The performed statistical tests for the equivalence of these groups lead to a rejection of their equivalence; nevertheless, one may argue that this rejection is mainly driven by the fact that the “waiting" subgroup exhibits a much stronger tie to the labour force than the unemployed. Within the marginally-attached population, we also observe that those individuals who have searched for a job but are still classified as inactive[30] display a transition rate into employment which is comparable to the long-term unemployed, even after controlling for individual characteristics. In addition, we find substantial heterogeneity among the non-attached (those inactives who do not want to work). This is due to the fact that the so-called long-term future job starters display the highest degree of attachment to the labour market, judging by its high average transition rate into employment (36.4%). Although these individuals also frequently move to non-attachment (33.3%), their transition behaviour is closer to unemployment than to the rest of non-employed. Therefore, its classification as inactive might not be adequate.

Overall, these results suggest possible shortcomings in those analyses which use slack measures based exclusively on the job search criterion, as it is the case of the unemployment rate. A broader analysis of the labour market seems to be appropriate for an accurate assessment of the labour market slack. Nevertheless, the results indicate that the job search and the reported desire to work provide meaningful and complementary information regarding the attachment of individuals to the labour market.

29. In addition, such differences are also reflected in the reported entry wages in each of these groups.

30. Either because they have searched passively or do not fulfil the other requirements for unemployment classification.

References

Brandolini, Andrea, Piero Cipollone, and Eliana Viviano (2006). "Does the ILO Definition Capture All Unemployment?" *Journal of the European Economic Association*, 4(1), 153–179.

Centeno, Mário and Pedro Fernandes (2004). "Labour Market Heterogeneity: Distinguishing between Unemployment and Inactivity." *Economic Bulletin Banco de Portugal*.

Centeno, Mário, Álvaro Novo, and José R. Maria (2010). "How to Measure Unemployment: Implications for the NAIRU." *Economic Bulletin Banco de Portugal*.

Clark, Kim and Lawrence Summers (1978). "The Dynamics of Youth Unemployment." *Working Papers, National Bureau of Economic Research*.

Clark, Kim and Lawrence Summers (1979). "Labor Market Dynamics and Unemployment: A Reconsideration." *Working Papers, National Bureau of Economic Research*.

Eriksson, Stefan and Dan-Olof Rooth (2014). "Do Employers Use Unemployment as a Sorting Criterion When Hiring? Evidence from a Field Experiment." *American Economic Review*, 104(3), 1014–39.

Flinn, Christopher and James Heckman (1983). "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?" *Journal of Labor Economics*, 1(1), 28–42.

Gönül, Füsun (1992). "New Evidence on Whether Unemployment and Out of the Labor Force are Distinct States." *Journal of Human Resources*, 27(2), 329–361.

Hausman, Jerry and Daniel McFadden (1984). "Specification Tests for the Multinomial Logit Model." *Econometrica*, 52(5), 1219–1240.

Hornstein, Andreas (2012). "Accounting for Unemployment: The Long and Short of It." *Working Papers, Federal Reserve Bank of Richmond*.

INE (2015). "INE - Instituto Nacional de Estatística, Documento Metodológico do Inquérito ao Emprego: Versão 3.1."

Jones, Stephen and Craig Riddell (1999). "The Measurement of Unemployment: An Empirical Approach." *Econometrica*, 67(1), 147–162.

Jones, Stephen and Craig Riddell (2006). "Unemployment and Nonemployment: Heterogeneities in Labor Market States." *The Review of Economics and Statistics*, 88(2), 314–323.

Kroft, Kory, Fabian Lange, and Matthew Notowidigdo (2012). "Duration Dependence and Labor Market Conditions: Theory and Evidence from a Field Experiment." *Working Papers, National Bureau of Economic Research*.

Krueger, Alan, Judd Cramer, and David Cho (2014). "Are the Long-term Unemployed on the Margins of the Labour Market?" *Working Papers, Brookings Papers on Economic Activity*.

Luce, Duncan (1959). "Individual Choice Behavior: A Theoretical Analysis." *New York: Willey*.

Martins, Fernando and Domingos Seward (2019). "Into the Heterogeneities in the Portuguese Labour Market: An Empirical Assessment." *Working Papers, Banco de Portugal, Economics and Research Department*.

Neves, David (2014). "Labor Market Flows After the Redesign of the Portuguese Labor Force Survey." *Master's Thesis. ISEG - Lisbon School of Economics and Management, Universidade de Lisboa*.

OECD (2018). "OECD Employment Outlook 2018." *OECD - Organization for Economic Cooperation and Development Publishing*.

Rua, António (2017). "Dating the Portuguese Business Cycle." *Economics and Research Department Articles, Banco de Portugal*.

Schweitzer, Mark (2003). "Ready, Willing, and Able? Measuring Labour Availability in the UK." *Working Papers, Bank of England*.

Tano, Doki (1991). "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?: New Evidence from the Gross Change Data." *Economics Letters*, 36(1), 113 – 117.

Yellen, Janet (2014). "Labor Market Dynamics and Monetary Policy." *Opening Remarks at the Federal Reserve Bank of Kansas City Economic Symposium, Wyoming, 22 August*.

Economics Synopsis
Non-technical summary

April 2020

The sources of wage variability in Portugal: a binge reading survey

Pedro Portugal

The availability of repeated observations on the wages of the workers, across different firms over a long period of time, enable the investigator to explore relevant information about *who the worker is* and *where he works*. When a worker establishes a match with a firm his wage will be partially determined by who he is, where he works, and by the success of the idiosyncratic match. This third component is called *match quality*. Besides the notion of who is the worker, where he works, and the quality of the match, arguably, there is an important dimension often neglected in the literature: *what he does*.

Another useful extension of this framework is to consider the human capital spillovers that may emerge from the relationship with other workers at the workplace. In this vein, the researcher wants to consider not only who the worker is, where he works, and what he does, but also *with whom*. The presence of peer effects may arrive from learning externalities or simply from peer pressure.

The allocation of workers to distinct firms and job titles is not random. Accounting for the endogenous *sorting* across firms and job titles is critical in the analysis of wage formation, namely, to investigate whether high and low wage workers match with high and low wage firms, to measure the importance of sorting, and to construct counterfactual exercises where workers are randomly assigned to firms and/or job titles.

In this survey, nine empirical studies exploring the sources of wage variability in the Portuguese labor market are summarized. The common feature of these studies is the attention given to worker, firm, and job-title heterogeneity. A number of important conclusions emerge from the studies.

First, worker heterogeneity is the most important source of wage variation contributing to one third of the wage variation. Firm effects were also found to be important contributing one fourth of the wage variability. For their part, job title fixed effects still explain an unsuspectingly non-negligible one fifth of wage variation.

Second, there is convincing evidence that the high wage workers sort themselves into high wage firms. Symmetrically, low wage workers tend to match predominantly with low wage firms.

Third, good (high productive) workers are more likely being to be teamed up with good (high productive firms) vindicating the notion of super-modularity or

positive assortative matching.The idea behind positive assortative matching is the complementarity between individual and plant productivity levels.

Fourth, there is convincing evidence that wage dispersion have remained constant for the past 20 years. In accordance with this evidence, it is shown that the contribution of the main components (worker, firm, job-title) in the variance of wages has been essentially constant over this period.

Fifth, real wages are quite sensitive to the business cycle. But, changes in the composition of the workforce may weaken this indication. Compositional bias, indeed, plays a very important role. Failure to account for worker, firm, and job title unobserved heterogeneity induces a countercyclical bias in wage cyclicality. This is due to the fact that low-skilled workers and low-productivity/ low-paying jobs account for a smaller share of employment in recessions than in expansions

Sixth, females tend to be less present in firms and job titles with more generous wage policies. Around one-fifth of the gender gap results from the non-random sorting of workers across firms (glass door), and another one-fifth results from non-random allocation of workers into job titles (glass ceiling).

Seventh, the ability of firms to set wages, arriving from frictions in the labor marker, are quite relevant. Monopsony power affects negatively the wages of workers, suggesting that firm market power is a key ingredient to explain firm heterogeneity in wage formation.

Eighth, after a firm closure, the post-displacement monthly wages are, on average, 7 percent lower than pre-displacement wages. This wage loss arrives roughly in equal parts from job degradation, loss of match quality, and negative sorting into firms with lower pay standards.

Ninth, the observed 17 percent union wage gap arrives in large part from the indication that unionized workers sort themselves into higher paying firms. In contrast, there is no evidence that union workers are systematically more productive than their non-union counterpart. A small part of the wage premium is engendered by elevated job titles and/or more generous promotion polices. Match quality, however, does not play a role in the determination of the union wage gap.

Tenth, schooling grants access to better paying firms and jobs. Around one fourth of the overall return on a year of education operates through the firm channel, whereas a third of the return operates through job-title channel. Peer quality has a sizable impact driving wages. In particular, one additional year of schooling of the co-workers generates, in average, a half a percent increase in the wage of the worker.

Economics Synopsis
The sources of wage variability in Portugal: a binge reading survey

Pedro Portugal
Banco de Portugal e NOVA-SBE

April 2020

Abstract
The availability of repeated observations on the wages of the workers, across different firms over a long period of time, enable the investigator to explore relevant information about *who the worker is* and *where he works*. When a worker establishes a match with a firm his wage will be partially determined by who he is, where he works, and by the success of the idiosyncratic match. This third component is called *match quality*. Besides the notion of who is the worker, where he works, and the quality of the match, arguably, there is an important dimension often neglected in the literature: *what he does*. In this survey, nine empirical studies exploring the sources of wage variability in the Portuguese labor market are summarized. The common feature of these studies is the attention given to worker, firm, and job-title heterogeneity.

1. The brave new world of worker, firm, match, and job title heterogeneity

In a perfectly competitive labor market, identical workers, working in homogeneous firms, should receive exactly the same wage, which would mean that the wages distribution would collapse to a single mass point. Even in the presence of search costs, the equilibrium outcome would still be unique, corresponding to the monopsony wage, a result known as the Diamond paradox, after Peter Diamond's (1971) celebrated paper. Observed wages, however, reveal an unsuspected degree of wage dispersion, even among observationally identical workers.

The theoretical construct of an equilibrium wage dispersion can, of course, be rationalized in a number of ways by considering heterogeneity across workers and/or firms either ex-ante or as the outcome of a monopsonistic wage competition (most notably under on-the job search, as in Burdett and Mortensen, 1998).

Acknowledgements: I am deeply thankful to Lucena Vieira for so patiently assisting me exploring the Quadros de Pessoal over the last 20 years. I am also grateful to António Antunes and Nuno Alves for valuable comments and suggestions. The analyses, opinions, and findings expressed in this article are those of the author and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.
Note: This article was submitted for approval on February 21, 2020, before the World Health Organization declared the COVID-19 a pandemic on March 11.

E-mail: pportugal@bportugal.pt

The empirical contribution of worker and firm heterogeneity to wage formation has been studied extensively, in particular after the seminal work of Abowd, Kramarz, and Margolis (AKM) (1999). By extending the conventional Mincerian wage equation to include worker and firm fixed effects, the study of AKM spurred a significant interest on the role of firms setting wages, uncovering the importance of heterogeneity of the firms' wage setting policies.

The access to matched employer-employee datasets, by enabling the researcher to follow the worker in the firm and his mobility over firms, allows the econometrician to measure and control for permanent worker and firm heterogeneity. In other words, by having repeated observations on the wages of the workers across different firms over a long period of time, the investigator grabs relevant information about who the worker is and where he works. In AKM models, this information is summarized in the worker and firm fixed effects. Worker effects are, of course, no more than averages of residuals, exhibiting whether or not the worker collects wages above or below what would be predicted by his observed characteristics. Similarly, firm fixed effects reflect the (assumed) constant generosity of the firms' wage policies.

What lies behind a worker fixed effect? Whatever respects the individual worker, is constant over the sample period, and influences the level of wages. Observed worker characteristics which do not change during the survey period such as gender, race, or birth cohort are implicitly accounted for in the wage regression with worker fixed effects. But permanent unobserved characteristics are also accounted for. Notoriously, the role of the workers' ability is subsumed in the worker fixed effect as is his family background or his motivation. Other features may also be unobserved to the econometrician but play a role driving the worker wage. For example, there is no shortage of studies showing that height, weight, or beauty have an effect on wages. The same can be said with respect to religion preferences or sexual orientation. The worker fixed effect catches all of those influences.

What lies behind a firm fixed effect? Whatever respects the firms' wage policy that persists over the survey period. Observed characteristics such as location and sector of activity are likely to affect wages and remain unchanged. Managerial ability, however, may well be unobserved. Economic rents, total factor productivity, and monopsony power may have an impact on wages but be unknown to the econometrician. Again, the firm fixed effect will condense all those effects.

Besides the notion of who is the worker and where he works, arguably, there is an important third dimension often neglected in the literature: what he does. This means that occupational or job title heterogeneity should also be considered in a wage regression, possibly by adding a third high-dimensional fixed effect. The presence of compensating wage differentials associated with the amenities or the hardship of the workplace would justify the inclusion of a job title fixed effect to the extent that it affects persistently the level of wages. Similarly, the risk of accidents or the health risks associated with specific occupations are also likely to engender compensating differentials. Union bargaining power may also be manifested at the job title level because collective agreements over wage floors are negotiated precisely at this level. Finally, the chronic overcrowding in some occupations will be translated into

persistently lower wages for those occupations. The job title fixed effect is expected to absorb all those observed and unobserved impacts.

Another useful extension of this framework, one that relies on the estimation of worker fixed effects and the identification of their co-workers, is to consider the human capital spillovers that may emerge from the relationship with other workers at the workplace. In this vein, the researcher wants to consider not only who the worker is, where he works, and what he does, but also with whom. The presence of peer effects may arrive from learning externalities or simply from peer pressure. In any case, the inclusion of the worker fixed effect in this extended wage regression circumvents the problem of homophily (workers sharing similar characteristics) and the inclusion of firm and job title fixed effects bypasses the problem of spurious correlation at the workplace level (Mansky, 1993; Angrist, 2014).

The allocation of workers to distinct firms and job title is not random. Indeed, workers search for employers and choose occupations maximizing the expected future stream of labor earnings (Roy, 1951). The inclusion of firm and job title fixed effects in the wage regression allows the econometrician to control for the endogenous sorting across firms and job titles, to measure the importance of sorting, and to construct counterfactual exercises where workers are randomly assigned to firms and/or job titles.

The estimation of worker and firm fixed effects enables the researcher to investigate how high and low wage workers match with high and low wage firms. The correlation coefficient between the worker and the firm fixed effects, if positive, will provide evidence that low wage workers are matched with low wage firms and that high wage workers populate disproportionately high wage firms. Until recently, the evidence on the sign of wage sorting was blurred, not least because measurement error in one fixed effect is transmitted with the opposite sign to the other fixed effect, biasing negatively the correlation coefficient.

When a worker establishes a match with a firm his wage will be partially determined by his fixed effect (who he is), by the firm fixed effect (where he works) and by the success (or failure) of the idiosyncratic match. To this third component we shall call match quality. In general, the match quality component cannot be disentangled from the work and firm components. However, under the assumption of orthogonality between the match quality and the worker and the firm fixed effects, a lower bound estimate of the match quality can be computed (Woodcock, 2015). In some instances, considering the match quality component can be very revealing.

In principle, the indication of a positive correlation between worker and firm fixed effects should not be confused with the notion of positive assortative matching, where high productivity workers are assigned to high productivity firms. Whereas it may be argued that the worker fixed effect reflects to a large extent his underlying productivity, there is no reason to establish that high productivity firms are necessarily high paying firms. In the absence of direct measure of firm's productivity, the sign and magnitude of assortative matching would be hard to grasp.

The inclusion of worker, firm, and job title fixed effects in the wage regression equation enlarges dramatically the set of research questions that can be addressed. In particular, the extended model allows the researcher to investigate the role of observed

and, in particular, the unobserved characteristics of the worker, the firm, and the job title. A useful methodology to study the impact of enclosing (persistent) unobserved heterogeneity is the extension suggested by Cardoso et al. (2016) of the decomposition method proposed by Gelbach (2016) in order to account for high-dimensional fixed effects. With this methodology one can ask, for example, by how much the gender wage gap would be reduced if workers were randomly assigned to firms? Is it true that the union wage premium is driven by the fact that unionized workers are more productive? Which part of the wage loss of displaced workers is due to the loss of match quality in the pre-displacement job? Do low productivity workers lose jobs disproportionally during recessions? Do workers move from high paying job titles to low paying job titles during downturns? How much of the return to schooling arises from working with more educated workers?

In this survey, nine empirical studies applied to the Portuguese labor market will be considered, where the wage equation including the three (high-dimensional) sources of heterogeneity (worker, firm, and job-title) is estimated and where the questions listed above, and others, are addressed.[1]

2. The architecture of the Portuguese wage setting system

The Portuguese Constitution gives registered trade unions the exclusive right to bargain over wages. Collective bargaining occurs mainly at the industry level but, irrespective of the level, collective agreements on wages determine wage floors only. According to the Labor Code, these union negotiated wage agreements only formally bind the members of employers association and the workers of signatory trade union(s). Collective agreements are of three types. First, and most importantly, there are industry level agreements negotiated between one or more employers' associations and one or more unions, known as *Contratos Colectivos de Trabalho* (CCTs). Second, there are collective agreements signed by several employers that are not part of an employers' association and one or more trade unions, known as *Acordos Colectivos de Trabalho* (or ACTs), that are significant in the financial sector and utilities. Finally, there are firm-level agreements between an individual company and one or more unions, so-called *Acordos de Empresa* (or AEs).

More concretely, collective wage agreements specify wage floors for the base wage for a set of job titles (*categorias profissionais*). Job titles are occupational categories whose specificity is sufficient to justify to bargain over. The detailed classification of the job titles accounts for the complexity of the tasks, the hierarchical standing of the worker, and the stress of the working conditions. In a typical year, around 300 collective agreements are signed covering close to 30000 job titles. Note that this classification is much more granular than commonly used occupation disaggregation.

1. The computational difficulties surrounding the estimation of this type of models are discussed in Guimarães and Portugal (2010), in which an algorithm overcoming those difficulties is offered.

Despite the legal strictures, the impact of collective bargaining agreements reaches far beyond the signatory parties. The critical mechanism shaping the formation of wages has been the systematic extension, via so-called *Portarias de Extensão*, of industry-wide agreements by the Ministry of Employment, following a request from either or both of the parties to the agreement. This means that, at the firm level, the are no longer wage differences between union and non-union members. Similarly, there will be no distinction between members and non-members of the employer association in terms of the wage agreement. The upshot of this near automatic procedure is that the main channel of union strength (as indexed by density) can only be observed at the collective agreement (typically, at the industry) level. Altogether, both agreements and their subsequent extensions explain levels of collective bargaining coverage in the order of 90 percent of workers. Firms frequently pay more than the bargained wage. In practice, around 40 percent of workers receive base wages above the bargained wage. Moreover, on the top the base wage, workers may receive meal allowances, overtime compensation, shift pay, and productivity and non-productivity related bonuses.

The difference between total wages and bargained wages is commonly termed *wage cushion*. There is some evidence that the wage cushion is more relevant in low paying sectors (see, in particular, Cardoso and Portugal, 2005). The explanation seems to be that employers in low bargained wage settings (most likely, industries with low union bargaining power) have the ability to set wages above the minimum to suit their particular labor market conditions.

Apart from this disaggregated or informal minimum wage apparatus, the share of workers receiving the national minimum wage has risen dramatically in recent years (the incidence of minimum wage earners currently exceeds 20 percent). Collective agreements and therefore extension arrangements can only pay less than the national minimum wage to specified types of workers such as apprentices and handicapped persons.[2]

3. A unique matched employer/employee/collective agreement/job title dataset

Researchers of the Portuguese labor market have been blessed with the access to a uniquely rich dataset which, by construction, matches workers with establishments while registering the worker's job title and the corresponding collective wage agreement.

The Portuguese data employed in the surveyed studies come from a longitudinal matched employer-employee dataset known as the Tables of Personnel (or *Quadros de Pessoal*) for the years 1986 to 2009 (excepting 1990 and 2001) and from its virtually identical successor survey the Single Report (or *Relatório Único*) for the years 2010 to 2017. This unique dataset is administered by the Portuguese Ministry of Employment,

2. For a detailed review of recent developments in collective bargaining and extension arrangements in Portugal, see Addison et al., 2017.

and is taken from a mandatory annual survey of all firms with at least one wage earner in the reference month, March of each year until 1993, October thereafter. The survey covers various firm and establishment characteristics, as well as a set of characteristics of the workforce. Being compulsory, it does not suffer from the non-response problems that often plague standard household and firm surveys. Further, the survey covers all Portuguese wage earners, with the exceptions of the Public Administration sector and domestic servants.

More specifically, the dataset includes information on the establishment (identifier, location, industry, and employment), the firm (firm identifier, location, industry, legal form, ownership, year of formation, employment, sales, and capital), and its workers (social security identifier, gender, age, education, skills, occupation, employment status, professional level, seniority, earnings [base wage, seniority-related earnings, other regular and irregular benefits, and overtime pay], normal and overtime hours, time elapsed since last promotion, professional category and the corresponding classification in a collective agreement).

4. The sources of wage variation

Why do similar workers receive different remuneration and why do similar firms pay different wages? In a labor market operating under perfect competition, each worker should receive a wage that equals his or her marginal (revenue) product. Wage differentials should reflect differences in worker productivity rather than depend on job or employer attributes (other than those affecting worker utility such as dangerous working conditions that will in normal circumstances attract a compensating differential). In turn, worker productivity has a basis in competence, whether observed or not, typically 'acquired' through investments in human capital.

There is no shortage of models seeking additional or alternative explanations for wage variability, but it is now the characteristics of firms that assume pole position. Given the plethora of such treatments (e.g., implicit contract theory, principal-agent models, and efficiency wages), two approaches take front stage. The first has a basis in rent-sharing/insider-outsider considerations, while the second emphasises labor market frictions.

Rent-sharing models predict that wages depend on the employer's ability to pay. In particular, wages are predicted to have a positive correlation with firm profits, since firms may find it beneficial to share their gains with their workers. These models explain why wages depend not only on external labor market conditions but also on firm productivity, profits, degree of competition, and turnover costs, as well as the bargaining strength of workers. They also explain why the wages of workers from different groups of occupations, educational categories, and seniority tiers are higher in some firms or industries than in others.

The other explanation for wage differentials among workers with similar characteristics targeted here derives from the job search and matching literature and emphasises the role of labor market frictions in wage determination. Thus, the

equilibrium job search model of Burdett and Mortensen (1998) predicts that firms may have incentives to offer higher wages than their competitors in order to guarantee a low quit rate and attract a large number of workers in a market characterized by the existence of frictions – even in circumstances of homogeneous workers and firms *ex ante*. For their part, matching models that also take into account the existence of frictions in the labor market provide an explanation for wage dispersion. Differences in match productivity in those models explain gains with their workers and pay above the going rate.

Another source of wage variation is job title heterogeneity, which may influence wage rates for a variety of reasons. First, it is well-known that tasks involving risks of fatal or otherwise serious accidents are better paid than safe tasks. Thus, for example, one should expect significant compensating differentials attached to occupations such as bullfighters (included in our sample). Second, jobs that need to be executed under difficult or stressful conditions are also expected to be more highly remunerated than jobs performed in pleasant environments; for example, one should observe higher wages for individuals working on offshore oil platforms or in mines. Third, the complexity of some tasks may require heavy doses of specific training or unusually skilled workers. This would be a reason why, for example, brain surgeons and prima ballerinas earn higher wages. Fourth, some occupations are known to be chronically 'overcrowded', whereas others are thought to be in excess demand. Finally, the kind of technology in use may also foster unionization of the workplace and favor rent seeking. Production activities that imply the concentration of a large number of workers in a single plant (say, in auto or shipbuilding industries) facilitate industrial action, and thus improved wage conditions. To properly incorporate these and other such wage determinants, one needs a very detailed accounting of the kind of jobs being undertaken by workers.

The work by Torres et al. (2016) were the first to explore the relative importance of worker, firm, and job title heterogeneity. The authors reported that worker heterogeneity is the most important source of wage variation in Portugal (contributing to one third of the wage variation). The unobserved component plays a slightly less important role than the observed non-time varying characteristics of workers such as gender and education. Firm effects were also found to be important (contributing one fourth), due in roughly equal parts to the unobserved component and to observed non-time varying characteristics such as regional location, capital ownership, and industry. For their part, although less important than either of the corresponding worker or firm effects, job title fixed effects still explain an unsuspectingly important one fifth of wage variation.

5. Wage sorting in the Portuguese labor market

In the Portuguese labor market, there is convincing evidence that the high wage workers sort themselves into high wage firms (Torres et al., 2016). The evolution of wage sorting, however, is quite striking. Torres and co-authors show that the correlation between the worker and the firm fixed effect steadily declined from 0.31 in 1991 to 0.08 in 2013. This evolution is in sharp contrast with the evidence presented for Germany by Card et al. (2013), which point to a sharp increase in wage sorting over time.

A number of labor market trends may have contributed to the decrease of wage sorting over the last 20 years of the survey. First, the average size of Portuguese firms has declined noticeably; second, the fraction of minimum wage earners has increased noticeably; third, the proportion of employed women has increased; and, finally, the fraction of college graduates has increased markedly. However, when each of the compositional change is isolated, there is no indication that the reported trends lie behind the evolution of the correlation coefficients. Indeed, it is apparent that the decrease of wage sorting is observed within each category. In particular, the trend towards less positive wage sorting occurs primarily among high wage workers.

6. Assortative matching in the Portuguese labor market

The sorting of heterogeneous workers across firms has been the subject of not inconsiderable debate. The idea behind positive assortative matching is the complementarity between individual and plant productivity levels, with good workers being teamed up with good firms. The theoretical basis for such matching is provided by assignment models. The early assignment models, however, were rooted in competitive equilibrium, thereby disregarding establishment-specific components in the wage equation. With the introduction of frictions, more recent developments have ensured a sorting of workers across plants. At issue in these models is the nature of the equilibrium matching pattern since different matching models predict different patterns according to the assumptions of the model.

There is a general consensus that good workers (i.e. the more productive ones) tend to earn higher wages. Therefore, it is possible to rank workers' productivity based on the individual permanent component of their wages, namely the worker fixed effects estimated from wage equations. Similarly, good firms (i.e. more productive ones) tend to have higher profits. However, these firms may pay lower or higher wages due to the presence of non-monotonicities in the wage schedule. Indeed, high-productivity firms have better outside options than their low-productivity counterparts, which may exert downward pressure on their workers' wages. This can be particularly relevant for low-skilled workers who may end up being paid less than if they were working for less productive firms. Non-monotonicities in the wage schedule also mean that wages reflect the marginal contribution to the value that the firm generates and it can be either the more productive or the less productive firms that derive higher marginal benefit from employing a better worker (Eeckhout and Kircher, 2011).

The estimation results given in Torres et al. (2018), in which the labor elasticities of a production function with thousands of different labor inputs are estimated, conveyed a consistent story in favor of the super-modularity or positive assortative matching hypothesis. The similarity between the magnitude of the wage sorting and the assortative matching correlations and, revealingly, their evolution over time seem to suggest that the correlations between the estimates of person and firm fixed effects from the wage equation may after all provide a sensible approximation to the measure of assortative matching. Contrary evidence from similar studies may be tied to the short

temporal dimension of the panels used. As argued earlier, fixed sample and limited mobility biases seem the likely culprits.

7. The sources of wage dispersion

The estimation of a wage regression with high-dimensional fixed effects provides a natural framework to investigate the sources of wage dispersion, and its evolution through time. The decomposition of the elements of the variance of wages is informative regarding the forces behind the changes in wage inequality.

When the breakdown of variance of (log) wages on an annual basis is established, Portugal, et al. (2018) show that the contribution of the main components in the variance of wages is essentially constant over the period under review. Indeed, it is not possible to establish a clear trend for wage dispersion when we consider each wage component separately (the variances of the worker, firm, and job title fixed effects).

It is not surprising, then, that the conventional indicators of wage dispersion have remained constant for the past 20 years. As mentioned above, the correlation between the fixed effect of the worker and the fixed effect of the firm, however, has been weakening, which in turn has led to less wage inequality.

8. The sources of real wage cyclicality

The cyclical behavior of real wages has been the subject of many studies since the debate of Keynes, Dunlop, and Tarshis. Earlier studies based on aggregate data showed some ambiguous results. One reason why these studies have reached no definitive conclusions resides in the fact that they have ignored the changes in the composition of the workforce over the cycle. Furthermore, with aggregate data it is implicitly assumed that the relationship between real wages and the business cycle is the same for all individuals or groups of individuals. If this assumption does not hold, then the estimates of real wage cyclicality are plagued by a specification bias.

The presence of compositional effects has attracted much attention in recent years and the seminal studies based on individual-level panel data for the United States show that cyclical changes in the composition of the workforce may generate a bias in the aggregate real wage.

It is also widely agreed that industry composition may change over the cycle. As pointed out by Okun, if some industries/firms offer rents to workers and if these sectors are also more cyclically sensitive, workers can switch into high-paying jobs during booms because such jobs are less tightly rationed during these times. This inter-industry/firm mobility of workers generates, per se, a procyclical behavior of wages.

After controlling for worker and firm heterogeneity, it can always be argued that the composition and quality of jobs within a firm are also likely to vary over the business cycle. If firms' promotions and hiring standards exhibit a cyclical pattern, overall wage cyclicality is mainly driven by workers changing job titles rather than by wage changes within job titles. The same line of reasoning applies if match quality among new hires

falls in a recession or firms hire proportionally more workers into low-skilled jobs in a recession than in an expansion.

In the study of Carneiro et al. (2012) it is shown that, accounting simultaneously for worker, firm, and job title heterogeneity, the data support the hypothesis that real wages in Portugal are quite sensitive to the business cycle, regardless of the measure of the cycle used. Compositional bias, however, plays a very important role. Failure to account for worker, firm, and job title unobserved heterogeneity induces a countercyclical bias in wage cyclicality. This is due to the fact that low-skilled workers and low-productivity/ low-paying jobs account for a smaller share of employment in recessions than in expansions

An important finding presented by Carneiro et al. (2012) was that entry wages are more responsive to the cycle than wages of existing workers within the same firm-job title. This finding is valid whether we use the unemployment rate as the measure of the business cycle fluctuations or the most common measure in search and matching models, the aggregate labor productivity. It is also shown that, after account for worker, firm, and job title heterogeneity, wages for all types of workers exhibit a wage-productivity elasticity close to unity. Finally, disentangling the job finding and job separation probabilities, real wages react positively to changes in the job finding probability and negatively to changes in the job separation probability.

9. The sources of gender discrimination

In a market where there are good and bad firms, or simply firms with different wage policies, segregation of workers across firms is likely to lead to a wage gap that will persist over time (Groshen, 1991).

Similarly, the allocation of workers to differently payed occupations or job titles, in particular, the restrictions to access those that are highly payed (glass ceilings), may also generate wage penalties.

Cardoso et al. (2016a) investigate the mechanisms that shape the gender wage gap in Portugal and provide a clear measure of the impact of the allocation of workers to firms and job titles. They found that one-fifth of the gender gap results from the segregation of workers across firms, and another one-fifth results from job title segregation.

Not only are women sorted more frequently into low-wage firms, but also the wage penalty increases with the size of the firm fixed effect. The authors conclude that the 'glass ceiling effect' operates mainly through worker allocation to firms rather than occupations.

Why are women penalised by allocation across firms? One possibility would be that female workers are less efficient job searchers. A number of mechanisms may lie behind this. First, women may search less intensively than men because they may have more productive alternatives for the allocation of their time (e.g., domestic production). Second, women may limit the set of job opportunities to jobs with particular characteristics (e.g., flexibility of work schedules). Third, women may shape their search strategy to the labour supply decisions of their husbands, as implied by

first generation labour supply models. For example, those women may severely restrict the geographical boundaries of their job search. Fourth, women may have significant disadvantages compared with men, exploiting their social networks to locate suitable job offers. Fifth, women may underestimate the relevant distribution of wage offers, either because they expect to be discriminated against by some firms, or because they misrepresent their true value to the firm and bargain wages less aggressively than men (Card et al., 2015). A telling indication that women have lower expectations regarding wages is given by evidence that they report lower reservation wages than men for similar observed characteristics (Addison et al., 2004).

Why are women allocated to lower-paying job titles? What is behind job title segregation? The presence of barriers to entry into high-paying job titles, driven either by the hiring decisions of the employers or by the requirements of the job, is certainly one of the mechanisms at work. Job promotion decisions biased against female workers may also be at play. Overcrowding of job titles highly preferred by female workers may drive their corresponding wages down. On the other hand, the access to some high-wage job titles that are controlled by closed shop trade unions are frequently male-dominated (e.g., longshoremen). The allocation into job titles is responsible for around one-fifth of the total gender gap. There is, however, no indication that the wage gap between men and women increases significantly along the job title fixed effect dimension.

Of concern is the indication of the study by Cardoso et al. (2016b) warning that the allocation of female workers to firms and job-titles did not improve over the last two decades. In fact it deteriorated somewhat. By 2013, females tend to be less present in firms and job titles with more generous wage policies, in comparison with 1991.

10. Monopsony power in the Portuguese labor market

The assumption of a single market wage that would cause all employees to instantaneously leave the firm after a one cent wage cut seems unrealistic. The empirical evidence suggests the presence of considerable wage dispersion among workers with similar characteristics and among similar firms. In particular, the firm effects estimated in wage regressions can be thought of as arising from distortions in the labor markets. Search frictions in the labor market such as imperfect information on alternative available jobs, moving and learning costs, firm specific human capital, reputation costs, exploitation of rents, and worker heterogenous preferences over non-wage job characteristics, are sources of labor market power, and help explain why firms have market power and why the labor supply curve faced by an individual firm is not perfectly elastic. These search frictions in the labor market may generate upward sloping labor supply curves to a particular firm.

In a standard wage setting model this means that firms possess some power to mark down their wages below the marginal revenue product. This is in line with the "new monopsony" literature popularized by Manning (2003), in which employers gain some market power derived from search frictions when setting wages. Monopsony is not understood in the traditional sense of a unique employer in the labor market,

but instead as synonymous with imperfect competition, monopsonistic competition, upward sloping labor supply curve to the firm, or finite labor supply elasticity. A particular firm may face an upward labor supply curve even if there is no concentration on the demand side of the market.

Félix and Portugal (2017) estimate a measure of labor supply elasticity to the firm directly from the production function, at a very granular level. Investigating the impact of the elasticity of labor supply to a particular firm and firm total factor productivity on individuals' earnings, the authors use the Gelbach exact decomposition to understand how firm's monopsony power is associated with the firm's wage setting policy. The results suggest that a one standard deviation increase in the labor supply elasticity increases wages by approximately 1.5 percents. This means that monopsony power affects negatively the wages of workers. Also, they find evidence that the elasticity of labor supply is mainly correlated with the firm effects as hypothesized in the labor economics literature. This suggests that firm market power is a key ingredient to explain heterogeneity in wage formation.

11. The sources of wage losses of displacement workers

There is scant empirical evidence on the economic mechanisms that generate the wage losses of displaced workers. The study by Raposo et al. (2019) offers a novel evaluation of the sources of wage losses incurred by workers displaced due to firm closure, bearing in mind that wages in the previous job are a function of a set of worker characteristics (for instance, gender, education, and experience) that are expected to yield, in general, the same return on the previous job and on the subsequent job, and a set of firm, job title, and match characteristics that do not necessarily yield the same return in subsequent jobs. Hence, if wages primarily reflect characteristics, then individual wages will be highly persistent and largely invariant to where individuals work, and potential losses due to displacement will be negligible. If, on the other hand, firm, job title, and match-specific heterogeneity are important, then the costs of displacement incurred by workers could be considerable.

A displacement event could lead to the loss of occupation-specific human capital due to the difficulty of finding a job that uses existing skills optimally or due to the depreciation of specific human capital during non-employment spells. Human capital has a decisive role during the early phase of the joblessness spell because larger human capital endowments are initially associated with greater job opportunities and higher opportunity costs of unemployment that necessarily erode with the progression of the unemployment spell.

Wage losses of displaced workers can be related to the firm, job title, and match quality that existed before and after displacement. The approach of Raposo et al. (2019) provides a unified framework that allows them to separately identify the components of the sources of the wage losses associated with the worker-firm pair into the contribution of worker, firm, and match quality.

In this framework it is important to distinguish a good worker in a good firm from a good worker-firm match (that is, a match with higher quality). In the event of a displacement, a loss occurs if a high-quality job match between the worker and the firm is dissolved. Furthermore, match-specific human capital accumulated over the course of the employment relationship is permanently destroyed when a job separation occurs. Its value is lost to both match participants and to the society as a whole.

Raposo et al. (2016) show that post-displacement monthly wages are, on average, 7.2 log points lower than pre-displacement wages. The results showed that sorting into job titles plays a very sizable role in explaining the losses experienced by workers displaced through firm closures, accounting for 37 percent of the total average monthly wage loss. The loss of match quality effects also plays a significant role, accounting for 32 percent of the total average monthly wage loss. The remaining 31 percent are attributed to the negative sorting of workers across firms with different pay standards.

12. The sources of the union wage gap

In a regime of near-universal collective bargaining coverage one may nevertheless discern sharp union wage gaps due to the heterogeneous influence of unions in covered settings. In the study of Addison et al. (2015) it is shown that the estimates of the union density wage gap for total monthly earnings top out at approximately 24 log points.

Employing the, by now familiar, regression approach which accounts for work, firm, job-title, match quality heterogeneity they show that the observed union wage gap would be very small if workers were randomly allocated to firms; that union workers are no more productive than their non-union counterparts; that part of the wage premium associated with union workers enjoying elevated job titles and/or benefiting from more generous promotion polices is important; and that the match quality does not play a role in the determination of the union wage gap.

The stochastic structure of estimated fixed effects can usefully be explored to understand the mechanism of wage inequality, via the use of quantile regressions. In other words, a decomposition similar to the one used for the mean can be employed to investigate the role of sorting in terms of worker skills, the generosity of firm's wage policy, and the paying conditions of different job titles.

Taking into account unobserved worker heterogeneity and sorting across firms with distinct wage policies and job titles that are differently paid, it is shown that union density contributes to less wage dispersion. But unobserved worker heterogeneity also plays an equalizing role under unionism. Unionized low productivity workers have higher wages than their non-unionized counterparts. In contrast, the allocation of unionized workers by job title makes for increasing inequality. The indication that the wage gap between unionized and non-unionized worker increases along the firm dimension is quite strong throughout the wage distribution. The contribution to wage inequality, however, is relatively mild.

13. The sources of human capital spillovers

While literature on the returns to schooling has made remarkable progress over the last 50 years the role of worker sorting across employers or jobs has been neglected. To the extent that workers with different levels of education are not randomly allocated to firms, and firms' pay standards are heterogeneous, they could be a key channel yielding returns to education. A similar argument could be built over jobs. There are remarkable wage contrasts across narrow occupations, possibly driven by differences in their degree of riskiness, the amount of specific training required, or the technology used. Provided that education can grant a "passport" to better paying jobs, part of the overall return on education would operate through a job title channel.

The study of Cardoso et al. (2018) explores the role of firm- and job-level heterogeneity shaping the returns to schooling, quantifying the impact of sorting of workers across firms and job titles on the returns to education. Analysis of the role of firm- and job-level heterogeneity structuring the returns to education begs another question: What if peers matter? Fundamentally, the quality of a firm will depend on the quality of its human resources. They progress to explicitly acknowledge that work within a firm is not undertaken in isolation, but with coworkers.

It is shown that indeed schooling grants access to better paying firms and jobs. It reveals that one fourth of the overall return on a year of education (7.9 log points) operates through the firm channel, whereas a third operates through the detailed job the worker performs. The worker component is responsible for 38 percent of the return to education.

In the second part of their analysis the authors consider the role of co-workers shaping individual wages. For this purpose, they extend the high-dimensional fixed effects framework to incorporate a measure of their peers' quality (via the average of the co-workers fixed effects). They show that peer quality has a sizeable impact driving wages. In their preferred specification, a 10 percent increase in the measure of peer quality leads to a wage increase of 2.1 log points. Very importantly, they conclude that one additional year of schooling of the co-workers would engender a half a percent increase in the wage of the worker.

14. The sources of future research

From the econometric specification point of view, the natural extension of high-dimensional fixed effect regression models is, of course, the generalisation towards varying regression coefficients, as in Gibbons et al. (2019) and Guimarães et al. (2017). In this vein, a promising empirical strategy is the transposition from linear to non-linear models, as in Félix et al. (2019). The development of theoretical models, in particular search and matching models, giving support to a valuable reduced-form interpretation of the parameters of AKM type models, as in Kline et al. (2019), is very welcome. Finally, the study of multiple and complex layers of social connections (e.g., parents, children, siblings, spouses, peers, classmates, neighbours, and friends) , as in Eliason et al. (2019)

is also seen as a recommendable way to enhance our knowledge about the influence of social networks on wages.

References

Abowd, John M., Francis Kramarz, and David Margolis (1999). High wage workers and high wage firms. *Econometrica*, 67(2): 251-333.

Addison, John T., Mário Centeno, and Pedro Portugal (2004). Reservation wages, search duration, and accepted wages in Europe. Bank of Portugal working paper 04-13.

Addison, John T., Pedro Portugal, and Hugo Vilares (2017). Unions and collective bargaining in the wake of the Great Recession: evidence from Portugal. *British Journal of Industrial Relations.*, 55(3): 551-576.

Addison, John T., Pedro Portugal, and Hugo Vilares (2015). The sources of the union wage gap: the role of worker, firm, match, and job title heterogeneity. Bank of Portugal working paper 15-12.

Angrist, Joshua (2014). The perils of peer effects. *Labour Economics*, 30: 98-108.

Burdett, Kenneth and Dale T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2): 257-273.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1): S13-S70.

Card, David, Ana Rute Cardoso, and Patrick Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics*, 131(2): 633-686.

Card, David, Jörg Heining, and Patrick Kline (2013). Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics*, 128(3): 967-1015.

Cardoso, Ana Rute, Paulo Guimarães, and Pedro Portugal (2016). What drives the gender wage gap? A look at the role of firm and job-title heterogeneity. *Oxford Economic Papers*, 68(2): 506-524.

Cardoso, Ana Rute, Paulo Guimarães, Pedro Portugal, and Pedro Raposo (2016). The sources of gender discrimination. *Bank of Portugal Economic Studies*. 2(2):47-66.

Cardoso, Ana Rute, Paulo Guimarães, Pedro Portugal, and Hugo Reis (2018). The returns to schooling unveiled. Bank of Portugal working paper 18-05.

Cardoso, Ana Rute and Pedro Portugal (2005). Contractual wages and the wage cushion under different bargaining settings. *Journal of Labor Economics*, 23(4): 875-902.

Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal (2012). Real wages and the business cycle: accounting for worker, firm, and job title heterogeneity. *American Economic Journal: Macroeconomics*, 4(2): 133-152.

Eckhout, Jan and Philipp Kircher (2011). Identifying sorting - in theory. *Review of Economic Studies*, 78(3): 872-906.

Eliason, Marcus, Lena Hensvik, Francis Kramarz, and Oskar Nordström Skans (2019). Social Connections and the Sorting of Workers to Firms. IZA discussion paper.

Diamond, Peter (1971). A model of price adjustment. *Journal of Economic Theory*, 3(2): 156-168.

Félix, Sónia and Pedro Portugal (2017). Labor market imperfections and the firm's wage setting policy. Bank of Portugal working paper 17-04.

Félix, Sónia, Paulo Guimarães, and Pedro Portugal (2019). Price stickiness: microeconometric evidence using hazard models with high-dimensional fixed effects. mimeo.

Gelbach, Jonah B. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, 34(2): 509-543.

Gibbons, Charles E., Juan Carlos Suarez Serrato, and Michael B. Urbancic (2019). Broken or fixed effects? *Journal of Econometric Methods*, forthcoming.

Groshen, Erica L. (1991). The structure of the female/male wage differential: Is it who you are, what you do, or where you work? *Journal of Human Resources*, 26(3): 457-472.

Guimarães, Paulo, Fernando Martins, and Pedro Portugal (2017). Upward nominal wage rigidity. Bank of Portugal working paper 17-02.

Guimarães, Paulo, and Pedro Portugal (2010). A Simple feasible procedure to fit models with high-dimensional fixed Effects. *Stata Journal*, 10(4): 628-649.

Kline, Patrick, Raffaele Saggio, Mikkel Sφlvsten (2019). It ain't Where you're from, it's where you're at. Communication presented at LEED, New York.

Manning, Alan (2003). Monopsony in motion: imperfect competition in labor markets. Princeton University Press.

Manksi, Charles (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3): 531-542.

Portugal, Pedro, Pedro Raposo, and Hugo Reis (2018). The distribution of wages and wage inequality. *Bank of Portugal Economic Studies.* 4(1):1-14.

Raposo, Pedro, Pedro Portugal, and Anabela Carneiro (2019). The sources of wage losses of displaced workers: the role of the reallocation of workers into firms, matches, and job titles. *Journal of Human Resources*, forthcoming.

Roy, A. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 96: 201-229.

Torres, Sónia, Pedro Portugal, John T. Addison, and Paulo Guimarães (2018). The sources of wage variation and the direction of assortative matching. *Labour Economics*, 54: 47-60.

Woodcock, Simon (2015). Match Effects. *Research in Economics*, 69(1): 100-121.

www.bportugal.pt