1

Banco de Portugal Economic Studies

Volume IV

BANCO DE PORTUGAL
EUROSYSTEM

1
volume IV

Banco de Portugal
Economic Studies

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

Lisbon, 2018 · www.bportugal.pt

BANCO DE PORTUGAL ECONOMIC STUDIES | Volume IV - n.º 1 | Lisbon 2018 • Banco de Portugal Av. Almirante Reis, 71 | 1150-012 Lisboa • www.bportugal.pt • Edition Economics and Research Department • Design Communication and Museum Directorate | Publishing and Image Unit • ISSN 2183-5217 (online)

Content

Editorial

January 2018

The first issue of Banco de Portugal Economic Studies for 2018 contains three essays.

The first article, by Pedro Portugal, Pedro Raposo and Hugo Reis is titled "The Distribution of Wages and Wage Inequality". The article consists of a structural analysis of the evolution of wages between 1988 and 2013 using data from the *Quadros de Pessoal* up to 2009 and the *Relatório Único* from 2010, on workers with ages between 18 and 65, full-time and not working in agriculture. These are longitudinal data on which workers and companies are tracked over time, with more than 40 million annual observations, corresponding to almost five million workers and more than 600 thousand employers.

The analysis is based on regression models relating the (logarithms of) wages to age, years of schooling, size of the employer, tenure and gender. Given the longitudinal nature of the data, the authors were able to use techniques estimating fixed effects per worker, per company and per job.

The results obtained are informative about the characteristics of the labor market in Portugal. A first set of results concerns the link between the characteristics of workers and firms and the wages. The second set of results uses the models estimated to decompose the wage changes in a part due to the change in the composition of the characteristics of the workers (such as the increases in schooling) and in a part that is due to structural changes, *i.e.*in the valuations of these characteristics (such as the increase in salary for each additional year of schooling).

In the first set of results, the study confirms the existence of systematic differences attributable to gender, with men having higher median wages by about 23 per cent. Furthermore, these differences are increasing in wage levels. The median salary schooling premium is substantial, about 7 percent additional salary per year of additional education. This premium is smaller in lower wages (below 5 percent in the 20th percentile of the wage distribution) and larger in higher wages (almost 9 percent in the 80th percentile). Experience, represented by age, positively impacts wages and, just as schooling, it has larger effects on higher wages. Finally, the results show that wages tend to increase with the size of the employer.

The second set of results includes the estimation of real wage growth between 1988 and 2013, with a real growth rate of 53 percent in median wages but with a lower value, 47 percent, in the 20th percentile and a higher value, 66 per cent, in the 80th percentile. When decomposing between compositional change effects and structural effects, the improvement in workers' qualifications accounted for 78 per cent of wage increases, with the remainder being due to structural changes in the determinants of salary. However, the study finds a reduction in the importance of investment in human capital in the form of on the job training, particularly in lower wages. A relevant result was only made possible by the estimation of fixed effects for each worker and for each employer. Unlike other European countries, there has been a decrease in the matching between companies and workers in Portugal. It was initially found that the more generous firms (after controlling for observables) hired the better paid workers (also after controlling for the variables observed), but this correlation has been decreasing, an evolution that by itself tends to decrease wage inequality.

The second article in this issue, by Diana Bonfim and David Pereira is entitled "GDP-Linked Bonds: Design, effects, and way forward". In this paper the authors study the effects of a potential use of public debt securities characterized by interest payments as a function of GDP growth in the issuing countries.

Since the sovereign debt crises of the 1980s and subsequent episodes of default or public debt restructuring, there has been a growing interest in ways to finance States that ensure better risk-sharing with lenders and help prevent sovereign debt crises through mechanisms that protect against the effects of recessions.

In this paper, the authors estimate the potential interest savings or costs, for euro area countries between 2000 and 2015, assuming that these countries would have been financed by issuing GDP-linked bonds. It was also assumed that in these bonds the coupon rate would change directly in line with the real GDP growth rate, but with a lower limit of zero. It is also assumed that the new coupon rate and the new interest charges would have no impact on other variables, such as the GDP, the deficit or the public debt totals. The results of the exercise show that euro area countries would have been able to pay less interest in 2008-2009 and 2012-2013. These savings would have been offset by higher interest charges in other years, notably 2000, 2006 and 2007. These results illustrate the counter-cyclical mechanism implicit in GDP-linked bonds. The sovereigns would have paid less interest during recessions, but in return they would have paid more in periods of economic growth.

To quantify explicitly how much additional space countries would have for anti-cyclical fiscal measures, the paper presents another simulation exercise, assuming that the total deficit/surplus, the path of public debt, and economic growth would be identical to those actually observed. Thus, the savings or interest costs associated with the issuance of bonds linked to GDP would have a direct impact on the primary balance. Using the new interest payments, the authors simulated what would be the counterfactual primary balances. The correlation coefficients between primary balances and real GDP growth are interesting because a positive correlation between the two variables can be interpreted as indicating the existence of space for a government to take counter-cyclical fiscal measures. The results show that the correlation between primary balance and GDP growth would be significantly higher if GDP-linked bonds had been issued. In advanced countries the median correlation between GDP growth and the primary balance grows from 0.5 in the baseline scenario to 0.77 with the use of GDP-linked bonds.

Despite these results indicating potential gains from the use of GDP-linked bonds, in practice there may be problems of operationalization and specification of contracts that should be prevented. There is also the possibility of moral hazard problems, with incentives to increase borrowing amounts, and of adverse selection, as countries anticipating a higher probability of having default problems become more likely to use these instruments.

In the third paper, by Nikolay Iskrev and entitled "Term premia dynamics in the US and Euro area: who is leading whom?", author Nikolay Iskrev studies the dynamic relationship between the term premium embedded in AAA-rated sovereign bond yields from the euro area and the term premium embedded in US government debt securities.

Longer-term interest rates can be viewed as risk-adjusted averages of expected future short-term interest rates. The term premium represents the compensation required by investors in long-term bonds because of the risk that short-term interest rates in the future may not evolve as expected.

It is well known that sovereign bond yields in advanced economies tend to move together. One of the objectives of the article was to establish if this movement is also verified in the term premia incorporated in the yields. To analyze this question, econometric models of the time structure for interest rates in the Euro area and the US were estimated. On the basis of these models, the yield rate is decomposed into, on the one hand, the expectations regarding the evolution of short-term interest rates and, on the other hand, the term premia. Then, the co-movement between levels and changes in term premia is measured with correlation coefficients. Focusing on the example of 10-year bonds, the results obtained show that term premia in the euro area and in the US have a similar evolution in the period under study, with a dominance of periods where correlation is positive and very strong, often exceeding 0.9. However, it may be more correct to compare not the levels, but the variations over time of the term premia. Even so, during most of the sample period, the correlation is still positive and relatively strong, but correlations weaken for bonds with maturities of four years or less.

The second objective of the article was to explore the evidence on a causal relationship between the two risk premia, that is, to what extent can one say that movements in the term premium in one economy determine the movements in the term premium of the other. To capture the direction of causality the author uses the classic concept of Granger causality and complements it with two more recent measures, transfer entropy and directional connectedness.

The results show that two of the measures – Granger causality and directional connectedness – indicate a stronger causal impact of variations in

US risk premia on the Euro area. Transfer entropy results show an inverse relationship, *i.e.*the Euro area has a larger impact. The causal relationship between the premia in the Euro area and in the US is more pronounced in the long term, but causality patterns vary over time. However, whatever the measure the level of causality of one area over the other is relatively weak. Globally, the causality between the Euro area and the US term premia is relatively weak. Given this evidence, a more plausible explanation for the strong co-movement is that there are global factors that affect term premia in both regions.

It remains to be seen whether any trends to reverse economic globalization and the corresponding integration of financial markets will lead to changes in these results.

The distribution of wages and wage inequality

Pedro Portugal
Banco de Portugal and NOVA SBE

Pedro S. Raposo
Católica Lisbon School of Business and Economics

Hugo Reis
Banco de Portugal and Católica Lisbon School of Business and Economics

January 2018

Abstract
Over 25 years (1988-2013) the composition of employment and the determinants of wages have changed notably in Portugal. In this essay, the individual records of Quadros de Pessoal/Relatório Único are used to identify the structural and compositional changes in the distribution of wages, in Portugal. The workers' education level was the variable that most decisively contributed to wage increases over this period. Aggregate compositional effects, influenced by changes in schooling , are seen to be more relevant than the aggregate of the structural effects, which are essentially determined by the secular productivity growth. The horizontal shift of the wage distribution over time did not, however, contaminate the wage inequality indicators, which have remained essentially constant. This indication, is largely due to the fact that worker skills heterogeneity, firm wage policies heterogeneity, and job title heterogeneity have remained surprisingly unchanged. The association between firms with generous wage policies and high-wage workers has weakened significantly over this period, contributing in muted way to lower wage dispersion. (JEL: J24, J31)

What we do

Over the past 25 years, more explicitly the period between 1988 and 2013, the Portuguese labour market underwent deep changes. These changes reflected, among other things, a marked improvement in schooling for those in employment, a growing proportion of women in work, and the ageing of the working population. This paper explores the wealth of information contained in the individual records of the Quadros de Pessoal/Relatório Único to characterise the changes taking place in wage distribution, with a distinction made between compositional changes, and structural changes in the factors that determine wages. Special emphasis will be given within this

Acknowledgements: The authors thank Lucena Vieira for very competent research assistance, and the comments by António Antunes and Nuno Alves.

E-mail: pportugal@bportugal.pt; pedro.raposo@ucp.pt; hfreis@bportugal.pt

framework, to move in the indicators of wage inequality, with specific focus on the different reasons underlying wage variations.

What has changed

Between 1988 and 2013, the average years of schooling for dependent employees went up dramatically. The baseline was embarrassingly low (six years of schooling) rising to ten years, a figure that was even so less than satisfactory (see Figure 1). This move reflects on-going changes in compulsory education and a growing feeling among families that formal education was a worthwhile investment.

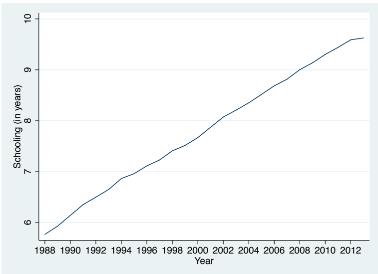

(a) Education

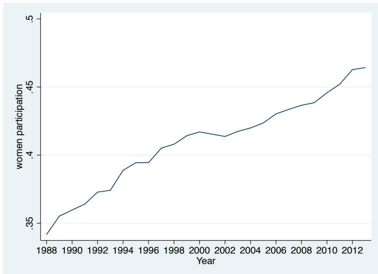

(b) Female Participation

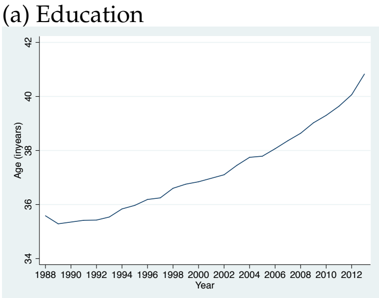

(c) Age

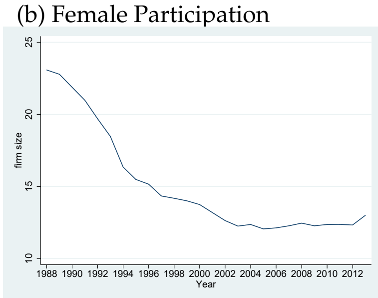

(d) Firm Size

FIGURE 1: Portuguese Labour Market Trends

At the same time, and likely due to the increase in educational levels, there was a growing participation of women in the labor market, with the female participation rate increasing from 34.2 per cent to 46.4 per in 2013 (Figure 1).

Two elements here — a rise in the age when youngsters left formal education, and an increase in retirement age — played a part in the rapid ageing of the active population, with the average age moving from 35.5 to 40.8 years (Figure 1). In spite of such a clear rise in the average age of workers, employment stability, measured by job tenure, remained unchanged during this period, standing at around 9 years.

One trend that does not seem to have been highlighted enough is the steep fall in the average size of Portuguese firms, with the average number of workers per firm falling from 23 to 13 during the period under review (see Figure 1). The fall in firm size came about fundamentally between 1988 and 2000, with large firms shedding labour and the number of micro-firms rising substantially.

Figure 2 shows the dynamics of base wage distribution through a breakdown into minimum, median and mean wages. The most relevant feature of this is, quite clearly, the horizontal shift in the distribution. This can be summed up in the real rises in mean wage (56.3 per cent) and in median wage (51.5 per cent). The growing density close to the minimum wage stems on the one hand from the elimination of wages paid at a rate below the minimum, affecting very young people in work durig 1987 and 1988 (Portugal and Cardoso (2006)) and, on the other hand, the growing incidence of the minimum wage influenced by nominal increases in the minimum wage above the aggregate figures for nominal wage increases (Martins and Portugal (2014)).

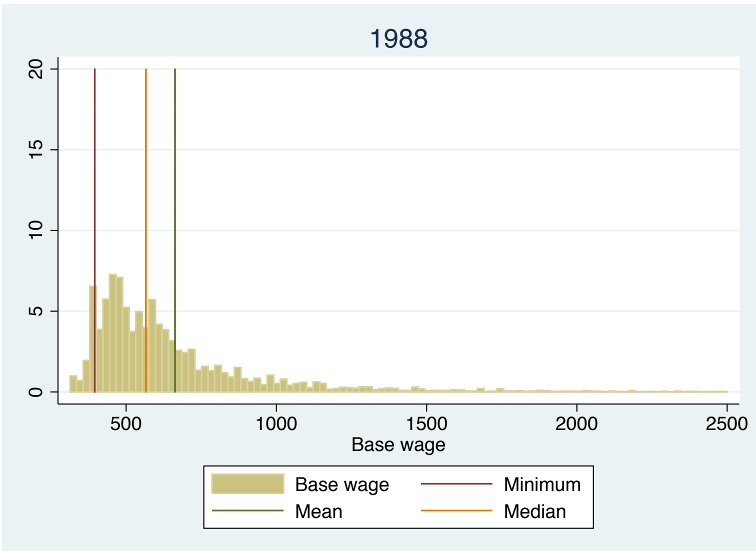

FIGURE 2: Evolution of the Wage Distribution (1988-2013)

Figure 3 also provides a dynamic indication of the shift in nominal wage change distribution for workers who hold a job in the same firm for two consecutive years (known as 'stayers'). The first point to make is the fact that nominal wage cuts are rare — indeed exceptional. This stems from the natural resistance among employers and workers alike to consider negative wage variation, but also, in a decisive way, of the prohibition (set down in the labour code) of any unilateral imposition of cuts in the base wage. The second point is that the combination of recession and low inflation leads directly to a pronounced increase in the fraction of wage freezes. This has, in recent years, come in at figures close to 70 per cent. These observations overall illustrate the well attested downward nominal wage rigidity of the Portuguese labour market (Martins and Portugal (2014); Nunes (2016); Addison *et al.* (2017)).

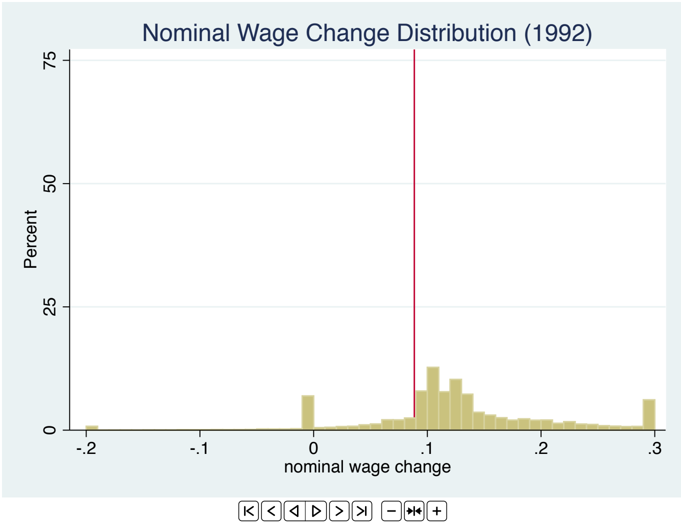

FIGURE 3: Nominal Wage Change Distribution

As a last point, summary indicators are presented for wage dispersion (Figure 4). The ratios between the total value of wages in relation to the 10th, 50th and 90th percentiles are surprisingly stable, contrary to international indicators. In fact, while on the left tail of the distribution, inequality in wages remained constant over the whole period under review, on the right tail there was the same uniform pattern, apart from the one-off increase in dispersion seen between 1989 and 1992.

In the following sections we will explore the factors contributing to changes in the wage distribution using quantile regression and high dimensional fixed effects regression to identify the sources of wage dispersion. First, however, before we move on to exploring these two lines of research, we will give a brief outline of the database that was used.

About the data

This paper has drawn on the microeconomic data gathered by the Ministry of Labour and Social Security, which brings together data from all establishements who employ at least one dependent employee: this is known as the Quadros de Pessoal (up to 2009) and then the Relatório Único (from

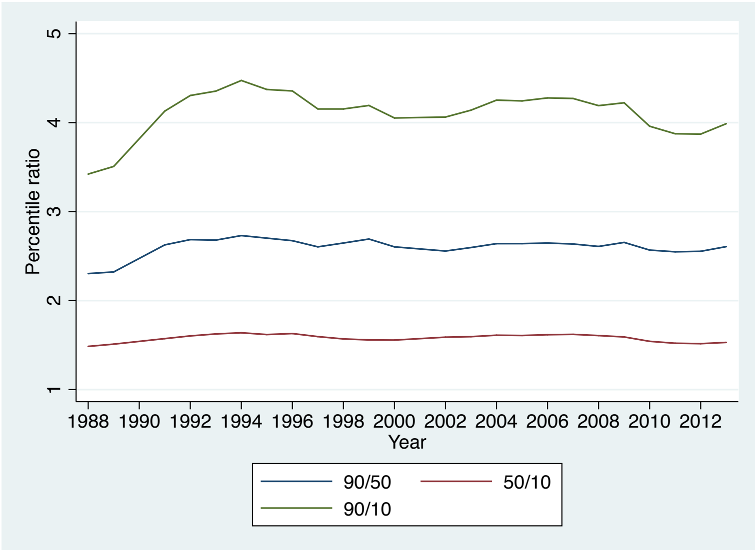

FIGURE 4: Wage Inequality

2010). The data obtained in these two reports is particularly precise and detailed when it comes to workers' wages. Initially the main purpose of the Quadros de Pessoal was to ensure that employers complied with wages reached through collective bargaining. In addition, it used to be a requirement for each employer to display the personnel tables in a public place. This was revoked in 2009, and a new directive allowed for the table to be shared digitally. Even today, in fact, the responsibility to "ensure that workers are provided with this information" remains in place.

Each individual worker is allocated a unique identifier, as well as each employer, every collective bargaining agreement and each distinct professional category. This allows us to track all workers systematically as they move along their career path. In addition to the detailed breakdown of wage components, the survey also collects social and demographic data on each worker, including age, gender, schooling, occupation, and job title. For firms, the survey collects data on the year the firm was founded, and includes sales revenue, size, location, and industry. This review will focus on full-time employees aged between 18 and 65, from all industries except the agricultural sector. Data is drawn therefore from 40,106,006 entries, relating to 4,918,285 employees and 611,765 firms.

Wage Setting

The usual way to analyse wage levels is to use an equation that accounts for levels of education and labour market experience. This equation has been adapted over the years from the now widely known equation initially derived by Mincer (1958). In this paper, we start off with an extended version of the Mincer equation to allow for differing wage levels according to gender (with a binary variable); the return on human capital specific to the firm; and wage level differentials linked to the size of the firm (through a logarithm of the number of employees).

The use of a linear regression model as a method to determine the factors linked to changes in wages levels has a number of drawbacks, however. To build a more complete picture, conventional regression analysis is not adequate to fully describe the wage distribution, since it is based on conditional means. In other words, in the regression model, the effect of a change in an explanatory variable, can always be shown as a horizontal shift of the distribution of a dependent variable. This statistical model is in fact unnecessarily restrictive, and in some cases inadequate, if the effect of an independent variable does not remain uniform across the distribution of wages. It is more useful to employ a more flexible model, which, in this case, is the quantile regression. This allows us to take into consideration the whole distribution rather than basing conclusions on the conditional mean.

Table 1 shows the results of estimating quantile regression corresponding to the 20th, 50th, and 80th percentiles for 1988 and for 2013. The indicator for gender discrimination is not homogenous across the wage distribution, with the gap more pronounced at higher wage levels. In this way, the difference between wages for men and women found on the 20th percentile is 18.2 log points (in 1998), which jumps to 24 log points when comparing wage quantiles corresponding to the 80th percentile. At the (conditional) median, men benefit by 20.9 log points.

One notable finding uncovered in the regression study is the indication that return on investment in formal education is significantly higher on the right tail of the distribution of wages. On percentile 20, an additional year of formal education leads to a rise in wages of 4.9 per cent, and at percentile 80 a rise of 8.3 per cent. This result suggests a strong complementarity between wages and education (Campos and Reis (2017)). The sequence of coefficients in the regression for the age variable, which is a proxy for the employee work experience, is similar to that seen for education. This also indicates that there is a complementarity between experience and more productive jobs. In terms of job tenure and firm size, the effects are relatively uniform across all distribution markers[1].

From 1988 to 2013, the main drivers influencing wages remained the same when it comes to education and length of service, but there is a drop in influence exerted by the firm size (across the whole distribution) and age

1. Another way to interpret this uniformity across distribution is to consider the fact that there is no heteroscedasticity (inconstant variance) associated with these variables.

Year	1988			2013		
Percentile	20	50	80	20	50	80
Male	0.182	0.209	0.240	0.170	0.251	0.339
Age	0.035	0.042	0.055	0.018	0.033	0.048
Age Squared	-0.000	-0.000	-0.001	-0.000	-0.000	-0.000
Tenure	0.001	0.001	0.001	0.001	0.002	0.001
Tenure Squared	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
Firm Size (logs)	0.068	0.073	0.075	0.037	0.049	0.054
Education	0.049	0.067	0.083	0.048	0.073	0.089
Constant	-1.737	-1.826	-1.961	-1.150	-1.636	-1.902

TABLE 1. Quantile Regression

(especially in the left tail) and the gender gap is more pronounced (particularly in the right tail). The change seen in the return on work experience is especially troubling as it may indicate lower levels of investment in job training for less qualified workers.

To better understand the changes in distribution of wages it is important to distinguish the effect of changes in the characteristics of workers and firms (composition effect) and the effect of changes in the regression coefficient (structural effect). In the linear regression models, this breakdown has most commonly been achieved using the Oaxaca-Blinder methodology. In the context of quantile regression models, the method proposed by Machado and Mata (2005) provides an elegant solution allowing for an aggregate breakdown between compositional and structural changes, and for a flexible way to identify the most decisive variables in the aggregate breakdown.

Table 2 shows the evolution of the compositional and structural effects for 1988 and 2013. It is very clear that compositional changes had a far more decisive impact on wage distribution than changes seen in the regression coefficients across the percentiles observed. In short, 33.3 log points (39.3 per cent) of the 42.7 log points (53.3 per cent) in the rise in median wage are due to changes in characteristics, and the remaining 9.4 log points (9.9 per cent) are due to changes in coefficients[2]. The significance of aggregated composition effects rises for higher percentiles.

	1988 (1)	2013 (2)	(2)-(1)	Composition Effect (4)	Structural Effect (5)
Quantile 20	-0.346	0.040	0.387	0.251	0.135
Quantile 50	-0.018	0.409	0.427	0.333	0.094
Quantile 80	0.420	0.925	0.505	0.409	0.096

TABLE 2. Machado and Mata Decomposition

In order to get a fuller picture of the drivers behind changes in the wage distribution, we need to consider a higher number of percentiles. Figure 5 is based on 99 regressions, corresponding to percentiles 1 to 99. This shows two wage distributions: what prevailed in 2013 and what would have prevailed if the characteristics had been the same as in 1988. In line with the results set out in Table 2, we can see a clear shift between the counterfactual distribution and the actual results, and this reflects an improvement in productive characteristics observed over 25 years. Figure 5 also presents two distributions: what prevailed in 1988 and what would have prevailed in 1998 with the same characteristics but this time applying coefficients from the regression survey of 2013. It is very clear that the change that occurred in relation to wage structures show definite improvements in wages, particularly for workers in lower income brackets.

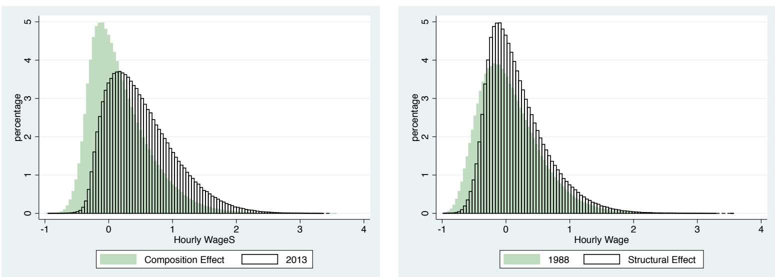

(a) Composition Effect (b) Structural Effect

FIGURE 5: Changes in the wage distribution - part I

2. The conversion of logarithm points into percentages is calculated using the generic formula $(exp(x) - 1) * 100$.

Finally, we looked at isolating the effects of higher education for workers and changes seen in the returns to labour market experience. Figure 6 compares the distribution in wages in 2013 with what would have been seen had the education levels matched those of 1988. The striking resemblance between Figure 5, which aggregates all compositional effects, and Figure 6 suggests that the additional years in formal education have been the most decisive factor in shaping the pattern of wage distribution. In fact, the increase in formal education is responsible for around 75 per cent of a rise in wages. Figure 6 shows the effect of the drop in the returns to labour market experience. Here we compare the distribution of wages seen in 1988 and the distribution of wages that would have been seen had the return on work experience been the same as in 2013. From this comparison we can see very clearly that the value placed on work experience for less qualified workers has dropped significantly, either because of a cut in investment in training or because the labour market no longer values these attributes in workers.

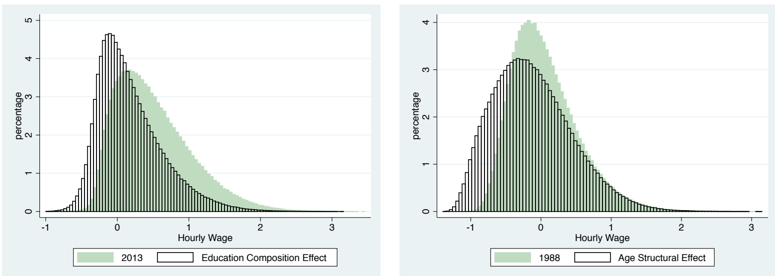

(a) Education Composition Effect (b) Age Structural Effect

FIGURE 6: Changes in the wage distribution - part II

On the factors underlying wage variation

The longitudinal data base (i.e. the records of workers spanning over various periods) makes it possible to raise questions which would not otherwise be feasible. In particular, the repetition of observations covering workers, firms and job titles will make it possible to obtain permanent effects. In essence, the sources of wage variation will be obtained through to three questions: Who is the worker? Where does he work? What does he do?

The worker dimension condenses the human capital which leads to higher or lower wages. The firm dimension summarizes the generosity of the firm wage policy. The job title dimension accounts for different remunerations that persist over time between occupations or tasks. Identification of these effects (known as fixed effects) is obtained from repeated observations of the worker, the firm and the job title and from the mobility of workers entrying and exiting firms and job titles. No particular interpretational problems are posed by the specification of a high dimensional fixed effect regression model (Guimarães and Portugal (2010)). All that is needed is to consider a Mincer regression extended to include 4918285 dummies identifying workers, 611,765 identifying firms and 127,021 identifying job titles[3].

In this context, heterogeneity of workers accounts for 49.2 per cent of the wage variance, whereas firm heterogeneity for 24 per cent and job title heterogeneity for 9.7 per cent[4]. This breakdown is similar to that obtained by Torres *et al.* (2013).

The correlation between the fixed effects of workers and of firms was estimated at 0.25, indicating that better paid workers (with greater human capital) tend to work in high paying firms (positive assortative matching). This association is fundamentally determined by the observed component of the worker (education, experience, tenure, and gender), where the correlation is 0.29, while the correlation with the unobserved part is 0.05.

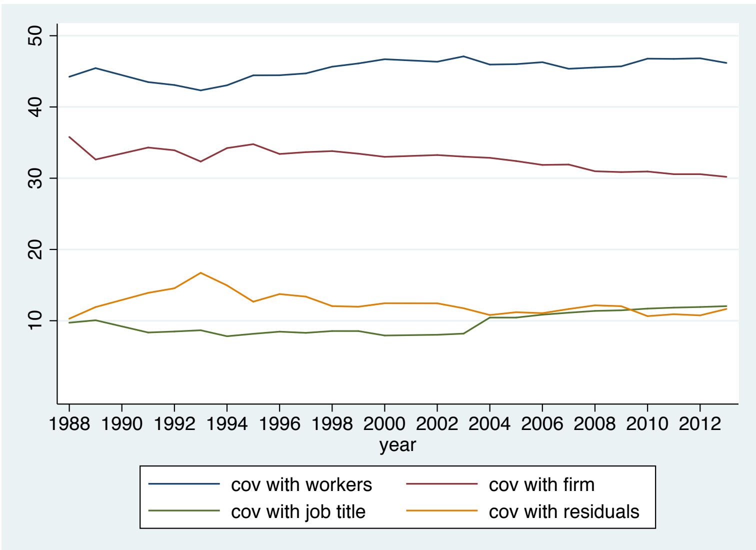

FIGURE 7: Wage Variance Decomposition

When the same breakdown of variance on an annual basis is established, it can be seen that the contribution of the main components in the variance of wages is basically constant over the period under review (Figure 7). The component with a marked trend is the correlation between workers and firms fixed effects. This drops by nearly 20 p.p., mainly from 1992 onwards (Figure

3. In practice, the regression model includes time-varying coefficients for gender, education and firm size as well as the three high dimensional fixed effects.

4. The expression used for this effect was $Var(Y = X + Z + W) = Cov(Y, X) + Cov(Y, Z) + Cov(Y, W)$.

8). It implies a fall in wage inequality, in marked contrast to the empirical evidence for Germany (Card *et al.* (2013)). It is clear that there is a significant weakening in the association of high paying firms and the level of human capital among workers (the evolution over time is given by Figure 9).

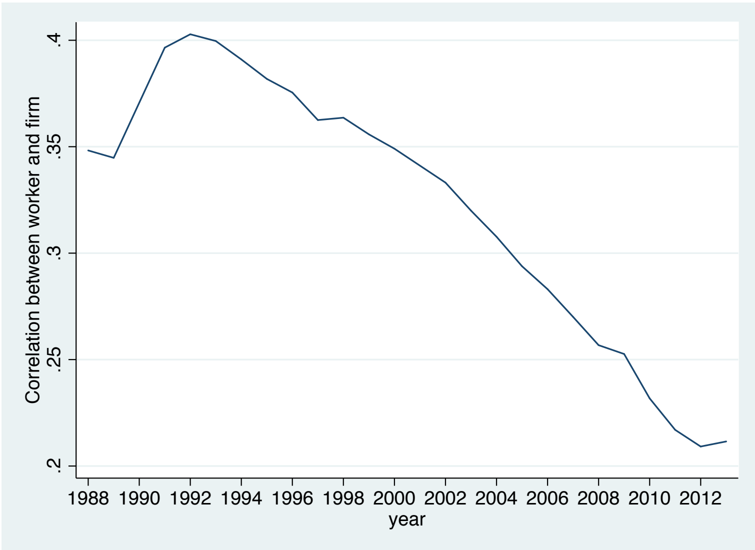

FIGURE 8: Correlation between Worker and Firm Fixed Effect

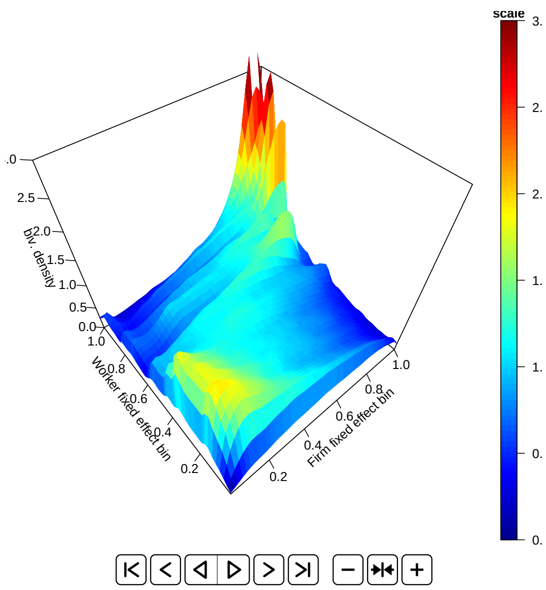

FIGURE 9: Matching between workers and firms

Summing Up

Over the 25 years from 1988 to 2013, the wage distribution of workers underwent profound changes. The most important of these was, as might be expected, the overall rise in wages, with the median wage increase reaching 53.3 per cent. Improvements in workers' qualifications accounted for 78 per cent of the increase, with structural changes in the determination of wages accounting for the remaining 22 per cent. The most decisive factor underlying the wage increase was a pronounced improvement in workers' education levels. According to our calculations, the raising of the school levels accounts for around three-quarters of the overall wage increase. Against this backdrop, there are indications that the second engine driving the production of human capital, on job training, slowed markedly during the period under review, mainly for those workers who are less qualified and collect lower wages.

Which of the questions below is the most important to best estimate individual's wages? Who are they? Where do they work? What do they do? The information on the worker is in fact the most relevant, since it explains around 50 per cent of the wage variation, though the information on the firm and the job is also important, since this explains a quarter and a tenth (for the firm and the job) of the wage variation.

It is not possible to establish a clear trend for wage dispersion when we consider each wage component separately. It is not surprising, then, that the indicators of wage dispersion have remained constant for the past 20 years. 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. This finding is in contradiction with most economic literature, and in fact indicates that the sorting of workers among firms is increasingly less determined by the complementarity between the worker's human capital and the productivity of the firm.

References

Addison, John, 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.

Campos, Maria and Hugo Reis (2017). "Uma Reavaliação do Retorno do Investimento em Educação na Economia Portuguesa." *Revista de Estudos Económicos Banco de Portugal*.

Card, David, Jörg Heining, and Patrick Kline (2013). "Workplace Heterogeneity and the Rise of West German Wage Inequality," Quarterly Journal of Economics." *The Quarterly Journal of Economics*, 128.

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

Machado, José A. F. and José Mata (2005). "Counterfactual decomposition of changes in wage distributions using quantile regression." *Journal of Applied Econometrics*, 20(4), 445–465.

Martins, Fernando and Pedro Portugal (2014). "Ajustamentos Salariais durante a Grande Recessão." *Boletim Económico do Banco de Portugal*.

Mincer, Jacob (1958). "Investment in Human Capital and Personal Income Distribution." *Journal of Political Economy*, 66, 281–302.

Nunes, André (2016). "Wage Adjustments under Extreme Downward Nominal Wage Rigidity." NOVA School of Business and Economics - UNL. Master Thesis.

Portugal, Pedro and Ana Rute Cardoso (2006). "Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations." *Journal of the European Economic Association*, 4(5), 988–1013.

Torres, Sónia, Pedro Portugal, John T. Addison, and Paulo Guimarães (2013). "The sources of wage variation: a three-way high-dimensional fixed effects regression model." Working papers, Banco de Portugal, Economics and Research Department.

GDP-linked bonds: design, effects, and way forward

Diana Bonfim
Banco de Portugal and
Católica Lisbon School of Business and Economics

David Pereira
Banco de Portugal

January 2018

Abstract
GDP-linked bonds have been proposed as a tool to help avoid sovereign defaults and debt restructurings. This article discusses potential advantages underlying the issuance of such an instrument, namely by quantifying the potential benefits that might arise when a country goes through periods of low growth rates and may face difficulties in meeting its financial commitments. The estimates suggest that there are potential benefits in terms of interest expenses. We simulate the correlation between primary balances and GDP growth in two scenarios: one with debt indexation to GDP growth and another one without such mechanism. As expected, the correlation between these two variables is significantly higher with indexation, suggesting that GDP-linked bonds could leave more room for automatic stabilizers to work during recessions. We run a similar exercise, but now considering a scenario where a country has to comply with a fiscal rule, and the main results are consistent. After establishing these facts, we examine recent issuances of GDP-linked bonds and discuss their limitations and weaknesses. This is crucial to understand what needs to be improved in the design of GDP-linked bonds to make them a universally used instrument. (JEL: E62, F34, H63)

Introduction

Sovereign debt restructurings have long been a concern both for investors and researchers. Most restructurings occur after sovereign default episodes and may have harmful consequences on the domestic economy and on the financial sector, leading to extended periods of exclusion from capital markets (Trebesch et al. 2012, Cruces and Trebesch 2013). While these concerns have been historically more focused on emerging market economies, the euro area sovereign debt crisis reignited this debate.

Acknowledgements: We would like to thank Nuno Alves, António Antunes, Vera Gouveia Barros, Cláudia Braz, Joaquim Cadete, Maria Manuel Campos, Isabel Horta Correia, José Miguel Costa, Jorge Correia da Cunha, Ana Filipa Fernandes, Miguel Gouveia, Filipe Morais, Pedro Moreira and Susana Salvado for their comments and suggestions. These are our views and do not necessarily reflect those of the Banco de Portugal or the Eurosystem. Any errors and omissions are our sole responsibility.

E-mail: dbonfim@bportugal.pt; dmtpereira@bportugal.pt

Against this background, discussions on the issuance of GDP-linked bonds have been recently in the spotlight.[1] The most prominent feature of a GDP-linked bond is the indexation of its coupon rate (or even its principal) to the issuer country's GDP growth rate, so that the security's cash flow payments would reflect the evolution of GDP.

This debt instrument could play an important role in helping to avoid solvency crises by, *inter alia*, providing countries with a form of insurance against downturns. It could, therefore, reduce the probability of defaults, debt restructurings and their associated costs. By avoiding high debt levels, it would decrease servicing costs and increase countries' fiscal space, allowing for less procyclical fiscal policies.

Researchers have been discussing GDP-indexed bonds since the 1990s and instruments with growth performance indexation features have already been issued (for instance by Costa Rica, Bulgaria, Bosnia Herzegovina, Singapore, and more recently by Argentina, Greece and Ukraine). However, this type of issuance is still considered an exception and it has not accomplished its full potential as an instrument that could play an important role in helping countries to avoid solvency crises and better risk sharing with private creditors.

This article begins by reviewing the existing literature on GDP-linked bonds. Then we describe the design of GDP-linked bonds, discussing how coupons could be determined. Afterwards we discuss the fiscal effects of GDP-linked bonds. We run three complementary exercises.

First, we try to quantify the potential fiscal benefits of issuing GDP-linked bonds, anchoring our estimates on previous work by Borensztein and Mauro (2004). To do so, we estimate the potential savings or expenses with interest for euro area countries between 2000 and 2015, assuming these countries had issued GDP-indexed bonds throughout the entire period. We also look separately into the potential effects for the countries at the core of the euro area sovereign debt crisis and for the other euro area countries.

Second, we estimate how much additional room countries would have had to pursue less procyclical fiscal policies. This is achieved by calculating, for the period between 2000 and 2015, the correlation between primary balances and GDP growth rates in two scenarios: with conventional bonds and introducing GDP-linked bonds. In the latter case, an "adjusted primary balance" is estimated considering the new interest amounts stemming from the introduction of the new instrument. We run these estimates separately for emerging market and advanced economies.

1. The G20, in the G20 Finance Ministers and Central Bank Governors Meeting of 24 July 2016, recognized that fiscal policy and fiscal strategies are essential in supporting growth. As such, G20 members, in that meeting's communiqué, called for "further analysis of the technicalities, opportunities, and challenges of state-contingent debt instruments, including GDP-linked bonds (...)" (G20 2016).

Finally, we run an exercise that resembles the previous one, but adding one additional feature. Specifically, we run our simulations considering boundaries on fiscal policies, such as those imposed by the Stability and Growth Pact. We run this exercise for France, Spain and Portugal.

All these partial equilibrium estimates rely on a strong set of necessary simplifying assumptions. All estimates should thus be considered as an upper bound on the potential benefits of GDP-linked bonds.

Before concluding, we summarize evidence on previous issuances of GDP-linked instruments and we discuss barriers to the implementation of such a product and possible solutions to overcome them.

Literature review

The international debt crisis in the 1980s led many governments to fail their legal obligation to meet debt repayments, in particular in Latin America and Eastern Europe. Ever since, there has been an interest in finding instruments that improve risk-sharing arrangements between governments and investors, allowing for a reduction of sovereign default probabilities and their corresponding costs. In this context, proposals for innovative financial products began to emerge, including the suggestion of indexing debt repayments to macroeconomic variables such as GDP, exports or commodity prices. As an example, Krugman (1988), in an attempt to solve the trade-off between debt forgiveness and financing, suggested that linking payments to measures of economic conditions could benefit both debtors and creditors.

Nevertheless, to understand better the importance of these types of instruments, it is crucial to describe the broader context in which they assume relevance. Debt restructurings are triggered by a default episode on debt commitments or by a debt-restructuring announcement. This is often when governments start negotiations with creditors, in order to agree on the terms of a debt exchange, providing debtors with debt relief. These debt restructuring processes are described as lengthy, costly and complex, most notably when compared to private sector processes (Bedford et al., 2005, Trebesch et al., 2012, Brooke et al., 2013). According to Forni et al. (2016), sovereign debt restructurings with external private creditors can, in fact, affect per capita GDP growth in the years after a restructuring.

In this context and given the frequency of financial crises, particularly in emerging economies, several authors have suggested ways to reduce inefficiencies of debt restructurings and their costs. For instance, Eichengreen (2003) discusses different approaches to this problem, presenting three main possible reforms: i) maintaining the *status quo*, while promoting the development of more complete and efficient debt agreements – a "contractual approach", falling under the currently used collective action clauses (CAC); ii) a "legislative approach" that would provide some of the functions of an international mechanism of insolvency; and finally iii) the establishment of a fully-fledged international bankruptcy court. The author suggests that those with reservations about these approaches would want alternatives, such as new forms of debt indexed to countries' real growth rate of their own GDP.

Research on GDP-linked bonds dates back to the early 1990s, with a proposal from Shiller (1993), who defends a market for long-term claims on the major aggregate income flows: GDP, occupational income, and service flows from commercial and residential real estate. Furthermore, he argues that instruments whose payments are linked to GDP could help reduce country risk and promote welfare.

One of the most studied advantages of GDP-indexed bonds is its ability to keep the debt/GDP ratio within a narrower range than conventional bonds. As a consequence, these bonds could play a role in preventing future debt crises, representing a way for countries to self-insure against possible growth downturns (Borensztein and Mauro, 2004). In this vein, Carnot and Summer (2017) investigate the reduction in uncertainty when issuing a fraction of public debt through GDP-linked bonds. Using the European Commission Debt Sustainability Monitor framework, the authors evaluate the impact on debt trajectories in the short and longer term. Their results indicate "important potential benefits" from the issuance of this instrument for all European countries, being especially relevant for economies with medium-to-high debt, high macroeconomic volatility and limited tools to smoothen shocks.

Cabrillac et al. (2017) estimate the possible gains for GDP-linked bonds´ issuers and conclude that the debt-to-GDP ratio would be reduced by 15% on average for a 25-year horizon for the 95th percentiles – the 5% least favorable simulated debt paths by 2040. They also defend that the volatility of the investors' portfolio would potentially decrease by 12% on average given the investment of such an instrument instead of investing in equities.

Interacting with the "debt/GDP ratio" effect is the fiscal margin that this product gives to countries to reduce the need to conduct pro-cyclical fiscal policies (Borensztein and Mauro, 2004, Blanchard et al., 2016). Borensztein and Mauro (2004) also investigate the particularly important benefits for countries that belong to economic monetary unions. Carnot and Summer (2017) consider the role of this instrument when monetary policy imperfectly responds to domestic shocks. Blanchard et al. (2016) argue that the introduction of GDP-linked could represent a "partial market-based solution to attain valuable insurance benefits" for euro area countries, ahead of a fiscal union.

Barr et al. (2014) develop a model of endogenous sovereign default, in which they analyze how GDP-linked bonds can raise the maximum sustainable debt level of a government and reduce the incidence of defaults. They use the concept of fiscal fatigue and standard debt dynamics equations to estimate debt limits, which will then be essential to model sovereign default with conventional and GDP-linked bonds. Under different risk aversion scenarios, the introduction of this security would increase the debt-limit level. In spite of this, investors demand a premium for providing insurance against GDP volatility. As the debt/GDP ratio increases, this specific cost gets overturned because the default premium increases accordingly.

There is also some research on the pricing of GDP-linked bonds. Borensztein and Mauro (2004) conclude that the insurance premium, that is the risk premium for holding bonds indexed to GDP to compensate investors for GDP volatility, would be small. Chamon and Mauro (2006) introduce the risk of default into their model. Firstly, they extract different combinations of probabilities of default and recovery rates from observed yields. Then, using the Monte Carlo framework, they simulate several paths for economic variables, including the debt/GDP ratio. Afterwards, they obtain a default trigger for the debt/GDP ratio and recovery rate that would yield the expected repayments implicit in the spreads. Finally, using the debt/GDP ratio default trigger and the simulated paths for the economic variables, they compute the corresponding payoff for both the growth-indexed bonds and the standard plain-vanilla bonds. The authors conclude that GDP-indexed debt can lower default frequency. When the share of this type of debt increases, both plain-vanilla and growth-linked bonds become less sensitive to GDP volatility and to growth shocks. Miyajima (2006) evaluates GDP-linked warrants (GLWs) considering the issuer's repayment capacity in the pricing formula. The author estimates the expected cash flows of debt payments, assuming that GDP follows a stochastic model, while trigger conditions are also modeled using the Monte Carlo framework. The issuer's capacity to service debt is defined as the difference between the incremental payments of GLWs and the increases in tax revenues due to economic growth. Finally, the author also uses the capital asset pricing model (CAPM) to calculate the size of the indexation premium, also finding it to be low.

Kamstra and Shiller (2009) estimate a risk premium of "only" 1.5 %. This estimate relies on the CAPM to calculate the cost of capital "relevant to issuing Trills" (a security with a coupon indexed to the United States' current dollar GDP, that would pay, for example, one trillionth of the GDP). They also defend that the maturity of this security would be long term, preferably perpetual. They consider that this new instrument would perform an important role as a stabilizer of the public budget.

Broadly speaking, GDP-indexed bonds, as put out by the International Monetary Fund (2017) when analyzing the economic case for state-contingent debt instruments for sovereigns, have the potential to "enhance policy space for sovereigns in bad states of the world, offer diversification opportunities to investors, and generate ancillary benefits for other economic agents and the broader system". Nevertheless, the institution recognizes some possible complications that, for some countries, may outweigh the benefits.

Sharma and Griffith-Jones (2006) also discuss the benefits of introducing GDP-linked bonds for borrowing countries, investors, the global economy and the financial system, while presenting the main concerns, issues and obstacles to their implementation. They also summarize recent experiences with these types of bonds, explaining their major flaws. Finally, in a similar vein to Borensztein and Mauro (2004) and IMF (2017), the authors defend the support of the official sector to help develop a specific market, suggesting several steps towards to this end.

Finally, sovereign equity-like instruments with some GDP-indexed features have already been issued. However, this type of issuance referred to warrants, attached to and often inseparable from an underlying bond and done in the context of debt restructurings. Benford et al. (2016) distinguish between potential GDP-linked bonds' issuances in normal and in debt restructuring times, with different benefits for issuers. During normal times they would help in preventing solvency crises, giving more fiscal space in downturns. In restructurings, this instrument would allow transferring higher debt repayments to periods when the economy is recovering.

The design of GDP-linked bonds: the coupon formula

The specific feature of a GDP-linked bond is the indexation of its coupon rate to the issuer country's GDP growth rate, so that the security's cash flow payments would reflect the evolution of GDP. In other words, the debt redemption's value would reflect the country's growth performance. If a government only issues this type of bond, all of its debt payments will change in line with growth. A GDP-linked bond coupon rate would equal:

$$coupon_t = \max(r + (g_t - \overline{g}); 0) \tag{1}$$

In order for the coupon rate to reflect the evolution of the GDP growth rate, an indexation factor, which would correspond to the difference between the observed growth rate (g_t) and a baseline growth rate ($\overline{g}$), would be added to the baseline coupon rate (r), thus linking coupon payments to economic performance.[2] This baseline growth rate, to be agreed at the moment of the contract, would reflect a trend growth rate and would adjust the economic performance of the year t to a period of growth of sufficient length. As such, if the economy in year t grows above the baseline growth rate, the indexation factor would be positive and the coupon rate would be higher than the baseline coupon rate. If the economy grows below the baseline, the indexation factor would be negative and the coupon rate would be lower than r. Finally, in order to protect investors from periods of particularly weak economic

2. For simplicity and in order to avoid another layer of risk, only the coupon rate – and not the principal – is adjusted. For further details, see for instance Borensztein and Mauro (2004).

performance – when, by adjusting the baseline coupon rate to a sufficiently negative indexation factor, the coupon rate would be negative – and thus also avoiding disincentives to investments in this kind of performance-linked security, a minimum of 0 would be applied to the coupon rate formula.

Therefore, the indexation of a bond to a country's economic performance would give governments a certain degree of insurance against periods of low growth rates. The magnitude of this protection would depend on the maturity of the bond and on the sensitiviness of coupon payments to growth. According to Barr et al. (2004), GDP-linked bonds with longer maturities offer sovereigns a better hedge against lower trend growth.

Fiscal effects of GDP-linked bonds

In order to quantify the insurance effect and to understand other potential fiscal benefits of GDP-linked bonds, a set of scenarios is presented, following some of the exercises laid out by Borensztein and Mauro (2004). We begin by estimating interest bill savings or expenses for euro area countries, should they have issued GDP-indexed bonds. Afterwards we run two exercises to estimate how much additional room countries would have had to pursue less procyclical fiscal policies: first we run a general exercise for advanced and emerging market economies; second we introduce fiscal constraints, running similar simulations but now assuming that there is a fiscal deficit limit of 3% of GDP (we run these estimations for France, Spain and Portugal).

We collect data from the IMF on the GDP real growth rate, on primary and overall balances as a percentage of GDP, on gross interest expenses as a percentage of GDP and on general government gross debt, also as a percentage of GDP.

Interest bill savings/expense

This first exercise is an attempt, through a simple approach, to illustrate how GDP-indexed bonds could affect a sovereign's interest bill. Following Borensztein and Mauro (2004), we consider a floating-rate bond with a coupon rate that follows a country's economic performance.

In this context, using equation 1, we simulate a new coupon rate and, accordingly, the amount of interest savings (or expenses) accumulated (or incurred). Underlying these simulations is the hypothesis that since the beginning of 1999 all the government debt of euro area countries consisted of GDP-linked bonds. It is also assumed that the new coupon rate and interest bill would have no impact on other variables, such as GDP, total deficit or debt, which, although unrealistic, could provide a measure of the expected potential amount of interest savings or expenses. Moreover, the baseline growth rate used corresponds to the average growth rate in the period 1992-2015, which should be long enough to provide a representative figure of the growth trend of a country.

As regards GDP growth, we use data in real terms, i.e. adjusted for inflation effects. It is true that indexing to GDP in nominal terms (as suggested by Benford et al. (2016)) would protect investors also from inflation fluctuations. However, it seems more prudent to spare both investors and issuers of another layer of complexity and risk, focusing solely on the countercyclical potential effect of real GDP-linked bonds.

As such, the actual implicit coupon rate is computed as the ratio of gross interest expenses of year *t* to the average of that same year's debt and the one of year *t-1*. However, it should be noted that this ratio does not consider that the actual debt stock also includes other instruments (such as currency and deposits and loans) and, more importantly, bonds issued in the past, under different market conditions. Also, one should take into account that countries that were under financial assistance were excluded from the bond market, leading to a less meaningful coupon rate. Finally, it is possible that GDP-linked bonds could have an additional risk premium, which is not considered in these estimates. All in all, these simplifying assumptions should make us cautious in interpreting the results, which should be regarded as the maximum savings awarded by this instrument.

The difference between each year's GDP growth rate and the baseline growth rate is added (or subtracted) to the coupon rate and the maximum of the adjusted coupon rate and 0 is computed. The new interest amount can thus be determined by applying the new coupon rate to the average of year *t* and year *t-1* debt.

In Figure 1 we present the results obtained for euro area countries, for the period between 2000 and 2015. The baseline growth rate of GDP considered in the exercises is 1.49%. Using equation 1, if euro area countries had issued GDP linked bonds throughout this period, they would have paid an average coupon rate of 4.34%. This is actually quite similar to the average coupon rates observed during this period (4.37%). As such, the aggregate savings on interest paid during this period would be negligible (0.13% of GDP).

However, these aggregate effects for the whole period hide important differences over time. Our estimates show that euro area countries would have been able to pay significantly less interest in 2008-2009 and in 2012-2013. This would have been compensated by higher interest expenses in several years, notably 2000, 2006 and 2007. This clearly illustrates the countercyclical mechanism embedded in GDP-linked bonds. Governments would have paid less interest in recessions, while paying more in periods of robust growth.

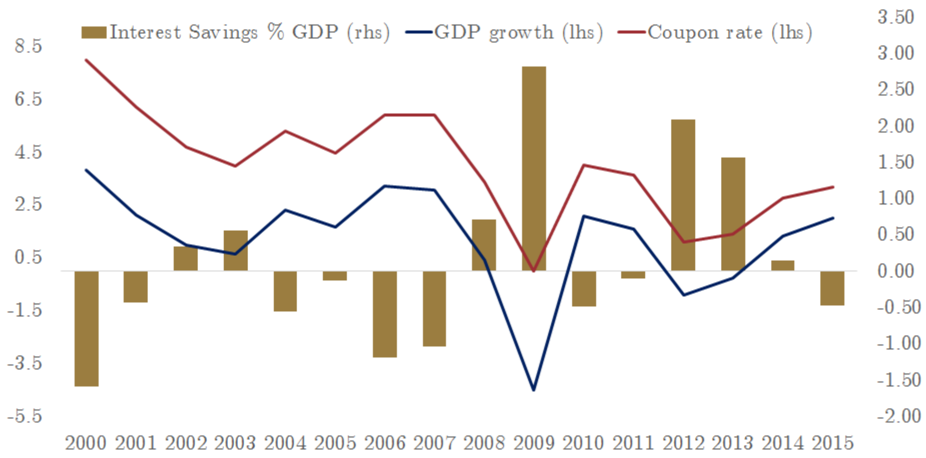

FIGURE 1: Interest savings as a % of GDP - Euro area

Source: IMF and authors' calculations.

Beyond the differences over time, it might also be interesting to consider differences between euro area member states. In Figures 2 and 3 we present the results of the same simulation exercise for two groups of euro area countries: those most affected by the sovereign debt crisis (Greece, Ireland, Italy, Portugal and Spain), and the remaining ones. The differences are striking.

For the crisis countries, the average coupon rate would have been 4.09%, significantly below the effective 4.35% observed during this period. This would entail savings with interest close to 0.3% of GDP. These savings would have been concentrated in the crisis years (2008-2014). In all the other years in the period under analysis, these countries would have paid more interest on their debt. [3]

For the remaining euro area countries, the pattern is much more irregular (Figure 3). There would have been interest savings in 2002-2003, 2008-2009 and 2012-2015. However, these are generally compensated by additional interest expenses in other years. The average coupon rate would have been 3.81%, only slightly below the observed average coupon of 3.91%. This would entail savings of 0.09% of GDP, i.e., one third of those potentially achieved by the crisis countries. These results suggest that GDP-linked bonds can generate interest savings even for advanced economies. However, given the caveats discussed above coming from the assumptions underlying this exercise (including the absence of a risk premium for these bonds), it is possible

3. We should note that the larger interest expenses for 2015 reflect to a large extent the strong economic recovery recorded by Ireland in this specific year.

that these benefits would be much smaller (or inexistent) in a more realistic scenario. As mentioned before, all these estimates are anchored on a set of simplifying assumptions that require some caution in their interpretation. To some extent, these numbers represent an upper bound to the potential interest savings achieved with GDP-linked bonds, for these countries, in this period.

One important assumption that can be relaxed is the inexistence of a risk premium attached to GDP-linked bonds (Benford et al., 2016). There is a lot of uncertainty on what this risk premium for euro area countries could be, most notably considering that at least initially there could be a novelty and a liquidity premium. We use the estimates provided by Kamstra and Shiller (2009) and re-estimate interest savings/expenses using a risk premium of 150 basis points. The average interest savings for the entire period would decrease from 0.13 to 0.08% of GDP in the euro area (from 0.30 to 0.22% in the GIIPS countries and from 0.09 to 0.05% of GDP in the other countries). Still, even with a 150 basis points risk premium, GDP-linked bonds could potentially yield interest savings across the board.

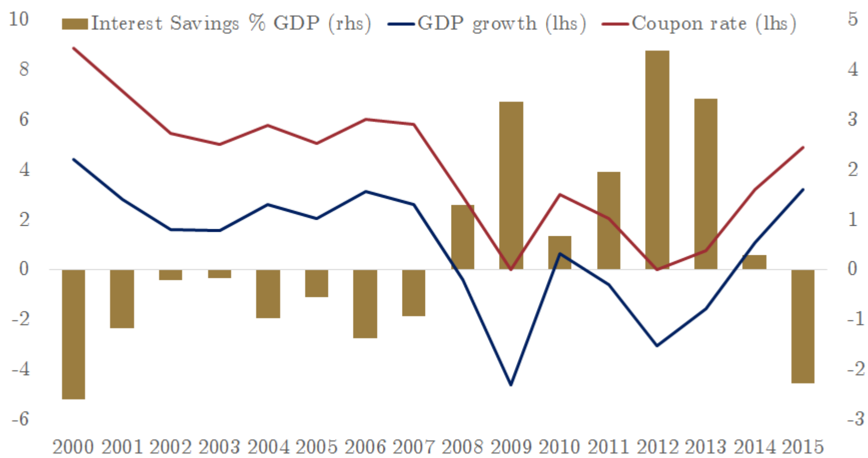

FIGURE 2: Interest savings as a % of GDP - Euro area crisis countries

Source: IMF and authors' calculations.

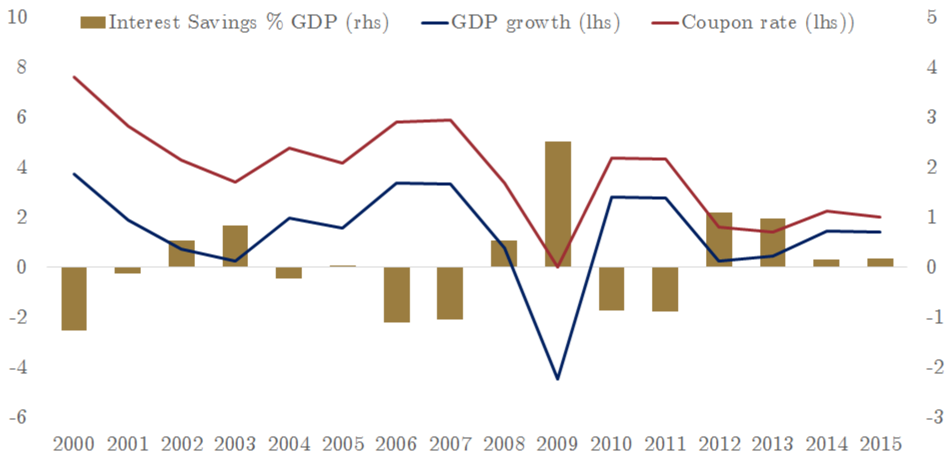

FIGURE 3: Interest savings as a % of GDP - Euro area non-crisis countries

Source: IMF and authors' calculations.

Generally, these results reinforce the conclusions of Borensztein and Mauro (2004), showing that when the GDP growth rate is below the baseline growth rate, the government generates interest savings with GDP-linked bonds. This would give room for pursuing policies that would result in a lower primary surplus (higher spending and/or lower taxes). It could also allow countries, in particular those that are following a short-term fiscal adjustment path, to achieve their fiscal goals faster. This would have been especially true for the countries at the core of the euro area sovereign debt crises, which underwent strong fiscal adjustments in order to regain market access. GDP-indexed bonds would thus provide countries with more fiscal space in times of crisis (allowing more room for the typical automatic stabilizers to work, without jeopardizing fiscal sustainability), while providing disciplinary mechanisms in times of growth (Brooke et al., 2013).

Fiscal policy

Mitigating procyclical fiscal measures. To better illustrate the countercyclical potential of GDP-linked bonds on fiscal policy, we replicate another exercise of Borensztein and Mauro (2004). The goal of this exercise is to explicitly quantify how much additional room would countries have had for countercyclical fiscal policy if their debt had been indexed to GDP. This is calculated by simulating the primary surplus that would have been obtained if all of a country's debt had been indexed to GDP growth. For that purpose, it was assumed that the total deficit/surplus, debt paths and economic growth would be the same as observed. It is thus assumed that, *ceteris paribus*, the interest savings or expenses stemming from the issuance of GDP-linked bonds would have a direct and proportional impact on the fiscal policy and thus on the primary balance. Other effects of a different fiscal policy, such as those relating to economic growth or risk premia, are not considered. These are of course very strong assumptions. While they are necessary to keep the simulations simple and tractable, they imply that these estimates are possibly not the same as those that would be obtained in a general equilibrium framework. We might thus interpret these estimates as an upper bound of the potential benefits of GDP-linked bonds.

We consider that in 1999 the entire debt stock had been indexed to GDP for 23 advanced countries and 15 emerging market countries.[4] The implicit interest rate is calculated as a ratio between the interest bill (taking gross interest expenses into consideration) and the average between the previous and the current year's debt stock. The "new interest rate" is simulated by applying equation (1) and adding the implicit interest rate to the "indexation factor", as previously described. The new interest amount is computed by multiplying that "new interest rate" by that year's debt. The baseline GDP growth rate corresponds to the geometric mean of the growth rates between 1980 and 2015.

The next step entails calculating the "adjusted primary balance", by using the new interest payments (maintaining the strong assumption that economic growth and fiscal variables are unaffected by the introduction of GDP-linked bonds). Finally, we compute the correlation between the simulated primary balance and the GDP growth rate. A positive and high correlation between these two variables can be interpreted as an indicator of a government's space to implement countercyclical fiscal policies. This correlation is compared to the correlation between the variables, but based on actual data. The results are reported in Table 1.

4. Advanced economies include Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Korea, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, United Kingdom, and the United States. Emerging market economies include Argentina, Brazil, Bulgaria, Chile, Colombia, Hungary, Indonesia, Latvia, Lithuania, Morocco, People's Republic of China, Peru, Poland, South Africa, and Turkey.

	Emerging markets			Advanced economies		
	Without indexation	With indexation		Without indexation	With indexation	
		Full indexation	Partial indexation (30% of debt stock)		Full indexation	Partial indexation (30% of debt stock)
	(1)	(2)	(3)	(4)	(5)	(6)
Mean	0.43	0.67	0.56	0.50	0.74	0.61
Median	0.51	0.75	0.67	0.50	0.77	0.59

TABLE 1. Correlation between the primary balance and real GDP growth, 2000-2015
Source: IMF and authors' calculations.

In fact, and in line with the conclusions of Borensztein and Mauro (2004) for a quite different period, in Table 1 we see that the correlation between the primary balance and GDP growth would be significantly higher with indexation than without it (comparing columns 1 and 2 for emerging markets and columns 4 and 5 for advanced economies).

To further enhance the realism of our estimates, we consider an alternative scenario, where instead of assuming that all the debt stock is composed of GDP-linked bonds, we consider that only 30% of the debt stock would be composed of this instrument. The results are also displayed in Table 1 (columns 3 and 6) and show that the correlation would still be higher than without indexation (though of course smaller than with full indexation).

This stabilization effect of GDP-linked bonds can be considered an automatic tool given their immediate and countercyclical fiscal reaction to growth - giving room for the typical automatic stabilizers to work freely during downturns and upturns. It can be argued that GDP-linked bonds offer a symmetric fiscal adjustment. They allow the channeling of fiscal revenues to interest expenses in good times, thus reducing the risk of overheating and at the same time relieving governments from the pressure of interest payments in bad times. [5]

5. According to the IMF (2015), fiscal stabilization reduces the volatility of growth over the business cycle. The institution estimates a potential decrease of around 20% of overall growth volatility for advanced economies, stemming from the move from average to high fiscal stabilization and a reduction of around 5% in the case of emerging market and developing countries. This is particularly important considering that higher fiscal stabilization and thus a lower level of growth volatility results in higher medium-term growth: "an average

Introducing fiscal constraints. We implement one final exercise, once more along the lines of the work by Borensztein and Mauro (2004). The aim of this exercise is to illustrate the ability to mitigate the effects of pro-cyclical fiscal policies by using GDP-indexed bonds for countries that belong to currency unions, such as the euro area, where the Stability and Growth Pact imposes boundaries on fiscal policy.

In this context, the exercise assumes that France, Spain and Portugal would have fully complied with the 3% of GDP limit on the fiscal deficit during the whole period. This is artificially achieved by imposing this limit each time that it was exceeded in actual data. The mechanics of the exercise are then quite similar to the previous one. We calculate the implicit interest rate as a result of the ratio of current year gross interest to the average of the previous and current year debt stocks. For simplicity, we assume that there is no feedback from the different deficit and debt levels on the interest rate or on growth.[6] A new debt path is computed following equation (2). This allows us to consider an adjusted primary balance that takes into account the 3% of GDP deficit limit.

$$\frac{D_t}{Y_t} = (1 + r - g_t)\left(\frac{D_{t-1}}{Y_{t-1}}\right) - S_t \tag{2}$$

In this equation, D_t refers to the debt stock, Y_t is GDP, and S_t is the primary balance as a share of GDP.

Following those same paths for debt and total deficits, a new primary balance is computed, but now supposing that all the debt stock was indexed to GDP growth. For the three countries considered in the exercise, we compute the correlation between primary balance and growth in a combination of two scenarios: (i) with and without GDP growth indexation; (ii) with and without the Stability Growth Pact limit. The four possible combinations of these two scenarios are reported for each country in Table 2.

strengthening of fiscal stabilization – that is, an increase in the fiscal stabilization measure by one standard deviation in the sample – could on average boost annual growth rates by 0.1 percentage points in emerging economies and 0.3 percentage points in advanced economies”.

6. Again, imposing these assumptions requires a cautious reading of the results. To fully capture all these effects, a general equilibrium approach would be necessary.

	France		Spain		Portugal	
	Without indexation	With indexation	Without indexation	With indexation	Without indexation	With indexation
Baseline estimates	0.63	0.82	0.92	0.96	0.17	0.66
With Stability and Growth Pact	0.51	0.87	0.78	0.90	-0.28	0.97

TABLE 2. Correlation between the primary balance and real GDP growth, 2000-2015
Source: IMF and authors' calculations.

When we compare the results with and without indexation, without imposing any limits on the deficit, the results obtained for France, Spain and Portugal are entirely consistent with those obtained for advanced and emerging economies in the previous exercise. The indexation of sovereign bonds to GDP significantly increases the correlation between primary balances and GDP growth. The largest increase is seen for Portugal, where this correlation is historically very low.

We can also gain some understanding about how fiscal boundaries within a currency union may limit the scope for countercyclical fiscal policy by comparing the results with and without the Stability and Growth Pact constraint. When we do so without indexing debt to GDP growth, we find that imposing a deficit limit of 3% of GDP would reduce a country's ability to conduct countercyclical fiscal policies, compared to the unrestricted baseline scenario. For France, applying this constraint would reduce the correlation between the primary balance and growth from 0.63 to 0.51, in the case of Spain from 0.92 to 0.78 and for Portugal, from 0.17 to -0.28. This is understandable, given that during downturns the possibility to increase the fiscal deficit (decreasing taxes and/or increasing expenditure) would be constrained.

Finally, we can quantify the benefits of indexation when the deficit constraint is active. We find that the correlation between primary balances and GDP growth is actually at its highest in this scenario for France and for Portugal (where this correlation actually reaches 0.97). However, for Spain, where the correlation is already quite high, there would be no apparent benefits from indexation in a scenario with fiscal constraints.[7] The benefits of GDP-linked bonds in terms of enhancing the space for countercyclical fiscal policies clearly depend on the starting point.

7. In the case of Spain, the indexation would entirely offset the procyclical effects imposed by the Stability and Growth Pact, according to our estimates.

It is important to note that all the exercises are anchored on assumptions that are necessary to conduct the simulations. However, these assumptions are specially strong in this third exercise, as both in the case of France and Portugal, the 3% GDP limit would have been binding for a large period of the sample (for Portugal it would have been biding throughout the entire sample period), making the comparison with the standard scenario more challenging.

Previous issuances, barriers to implementation and possible solutions

The introduction of GDP-linked bonds, as laid out in the previous sections, could be beneficial for borrowing countries. They could play an important role in avoiding solvency crises by, *inter alia*, increasing countries' fiscal space and allowing for countercyclical fiscal policies. As such, defaults, debt restructurings and their associated costs could be mitigated. Notwithstanding these advantages, the fact is that the issuance of instruments with these characteristics is considered an exception and has not been common on financial markets (Cabrillac et al. 2017). In this section we summarize evidence on previous issuances and discuss their shortcomings and barriers to implementation.

Previous issuances of sovereign's equity-like instruments

As the literature about equity-like instruments has been evolving, the issuance of this kind of products has also been somewhat progressing. In the end 1980s, as part of its debt relief within the "Brady Plan", Mexico pursued a debt-equity conversion program under which creditors (in this case, commercial banks) would be entitled to receive oil revenues owned by the country if its price exceeded a certain amount.[8] Also within the Brady Plan, other countries, such as Venezuela, Nigeria or Uruguay, have issued similar equity-like instruments. Later in the 1990s, and still part of the same plan, other countries such as Costa Rica and Bulgaria issued bonds for sovereign funding purposes, whose repayment was indexed to GDP, i.e. its payoff increased if GDP (or GDP per capita) of those issuing countries rose above a certain

8. The Brady Plan was announced in 1989 by US Secretary of Treasury, Nicholas Brady, in the context of the developing countries' debt crisis in the 1980s, which led some of them to default. As such, countries were settling rescheduling agreements with commercial banks, but without haircuts. The Plan, which was later (financially) supported by the IMF and the World Bank, consisted of debt reduction programs as a contribution to solving the above-mentioned crisis. The Brady Plan foresaw (i) exchange of outstanding bank loans into new sovereign bonds, partially collateralized by US Treasury bonds; (ii) a range of options of new instruments, such as discount bonds with a reduction in the face value, and par bonds with long maturities and below-market interest rates but no debt reduction and (iii) capitalization of interest in arrears to commercial banks into new short-term floating rates (Trebesch et al. 2012).

threshold. There are other examples of GDP-linked warrants' issuance, such as Bosnia Herzegovina and Singapore and more recently by Argentina, Greece and Ukraine.[9] The characteristics of some of those issues are summarized in the Appendix.

Overall, these issuances were mainly done in the context of debt restructurings, attached to (and often inseparable from) a conventional bond. Furthermore, their indexation formulas and conditions have usually been exceedingly complex, lacking standardization and clarity on the underlying reference data, as in the case of Bulgaria. In the case of Argentina, for example, as put out by Benford et al. (2016), the 350-day time lag between the reference (when payment is calculated) and payment date reduces the countercyclical effect and also suffers from great complexity. As such, despite all the apparent advantages of GDP-linked bonds described and quantified in this article, this instrument has rarely been used.

Barriers to implementation and possible solutions

There are important obstacles to the implementation and operationalization of GDP-linked bonds that explain why this instrument is not widely used, despite its conceptual advantages.

The main concern regards GDP data, in particular inaccuracies in its measurement and constant revisions (both due to revisions and updates in the underlying information and in methodologies), as discussed by Cecchetti and Schoenholtz (2017). The possibility of misreporting is also an important consideration. Indeed, data transparency and integrity is crucial from the investor point of view. In this context, increased independence of statistical agencies and technical support from international institutions could be decisive in guaranteeing the reliability of data, the accomplishment of statistical standards and in conveying credibility to investors. The risk of reporting manipulated data, however, seems somewhat contained by eventual reputational effects to the issuing sovereign. According to Borensztein and Mauro (2004), politicians' re-elections are supported by high growth rates, and thus it would not be reasonable to report, at least for several years, understated growth rates. Concerning data revisions, several authors suggest similar solutions to overcome this obstacle (Borensztein and Mauro, 2004, Sharma and Griffith-Jones, 2006, Brooke et al., 2013). The most important would be establishing ex ante (i.e. in the bond contract) the reference period for GDP data. Benford et al. (2016) suggest a six-month lag, but Cecchetti and Schoenholtz (2017) consider this period "inadequate". In any case, this lag

9. Portugal has issued a debt instrument called Treasury Certificates Savings Growth. This debt instrument is sold mainly to retail savers and part of its remuneration is indexed to GDP growth, thus having some features of a GDP-linked bond.

period should be long enough to have more accurate/precise estimations, but not so long so that the countercyclical effect would be lost.

Another obstacle is the absence of such market for these instruments and the corresponding concern with sufficient liquidity. According to Sharma and Griffith-Jones (2006), markets could be illiquid for this type of relatively new instrument. In this vein, and in line with these authors, there may be the need for a coordinated approach of several borrowers and institutions. This coordination should be both in terms of timing of issuance and conditions, following a standard design.

Moreover, it is arguable that such a new product, with an additional layer of complexity when compared to a plain vanilla bond, would be difficult to price, thus alienating investors. A possible way to overcome this obstacle could be by designing a simple and standard instrument, while having the technical support in pricing such a product.

Finally, as defended by Sharma and Griffith-Jones (2006), there could also be a moral hazard effect. Since higher GDP growth leads to higher interest payments, governments could have less incentive to implement policies to foster growth. This, however, as the risk of data manipulation, does not seem likely in the sense that lower growth would cost politicians both credibility and popularity. The IMF (2017) also mentions potential adverse selection problems coming from the fact that the countries which anticipate more negative macroeconomic scenarios might be the ones who are more eager to issue these instruments, thus raising their premia.

Given the solutions to the obstacles presented above, it is clear that international institutions could play a crucial role in overcoming them, namely by giving statistical support, monitoring data integrity, or using its published data as a reference. They could also help in designing a GDP-linked bond prototype, which could act as a standard model, and use its technical knowledge to enhance pricing. Their role, however, could be pushed even further. Sharma and Griffith-Jones (2006) argue that multilateral or regional development institutions could develop a portfolio of loans, whose repayments would be linked to the debtor country growth rate. These loans could be then securitized and sold on the international financial markets. International institutions already play an important role by giving financial assistance to countries. As such, when a country loses access to financial markets and needs financing from an international institution, this could be an opportunity for the country to sell to the institution GDP-linked bonds and for the institution to build the above mentioned loan portfolio. These financial assistance programs are accompanied by a reform package that, in principle, would increase potential growth. This does not preclude, however, the above-mentioned coordinated approach (in which these institutions could take a leading role). International institutions could also have a coordination role by, *inter alia*, guaranteeing that a sufficient volume of GDP-linked bonds is issued in order to reduce the liquidity premium (Cabrillac et al. 2017) and gathering a group of issuer countries that would allow to eliminate any potential reputational risk associated to countries with higher debt levels.

A recent noteworthy initiative to foster the GDP-linked bonds market is the London Term Sheet. This document describes in detail a template for the issuance of GDP-linked bonds, thus promoting the standardization of this product. This tool was developed by an ad hoc working group consisting of investment managers, lawyers from the private sector and economists from the Bank of England. This might provide the grounds for a standardized and transparent approach, with a direct involvement from the public sector.

Concluding remarks

Researchers have been discussing GDP-linked bonds since the 1990s and some sovereign equity-like instruments have been issued. However, this type of issuance is still considered an exception and has not, by far, accomplished its full potential as an instrument that could play an important role in helping countries to avoid solvency crises.

Theoretically, indexing a country's debt payments to its economic performance could give governments a certain degree of insurance against periods of low growth rates. As such, this article is an attempt to illustrate the potential advantages of the issuance of GDP-linked bonds, building up on previous work by Borensztein and Mauro (2004). Through three simulation exercises relying on a set of simplifying assumptions, we illustrated and quantified this insurance effect.

In a partial equilibrium analysis and relying on a set of assumptions, we show that the interest bill savings for the euro area countries at the core of the sovereign debt crisis could have been significant if they had issued GDP-linked bonds. These savings could have created room for less pro-cyclical fiscal policies, without jeopardizing fiscal sustainability. At the same time, interest bill expenses would have been higher during growth periods, thus contributing (albeit marginally) to promote a disciplining device to avoid excessive public spending during these periods.

Moreover, we find that the correlation between primary balance and real GDP growth is substantially larger when GDP linked bonds are used. This is true both for advanced economies and emerging markets.

These results should be read without forgetting the caveats and limitations of the simulation exercises conducted. For instance, the calculation of the (implicit) coupon rate (as a ratio of interests paid and the debt stock) does not take into account that the actual debt stock also includes other instruments (such as currency and deposits and loans). Furthermore, all the exercises hinge on the assumption that changes in some fiscal variables (e.g. interest amount) would not affect economic growth and other fiscal aggregates. Most of the exercises abstract from the existence of a risk premium that should be attached to these bonds. While these assumptions are necessary to keep calculations simple and intuitive, they necessarily imply caution in the interpretation of the results and of the policy implications.

That said, our results reinforce a vast literature illustrating the potential benefits of GDP-linked bonds. In this article we also discuss the main barriers to implementation and potential remedies to address them. These barriers are not unsurmountable and the recent interest from scholars and policymakers may offer the necessary solutions to widen the use of these instruments worldwide.

References

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

Barr, D., Bush, O., Pienkowski, A. (2014) "GDP-Linked Bonds and Sovereign Default." Bank of England, Working paper 484.

Bedford, P., Penalver, A., Salmon, C., (2005) "Resolving sovereign debt crises: the market-based approach and the role of the IMF." Financial Stability Review, 91–100.

Benford, J., Best, T., Joy, M. (2016) "Sovereign GDP-linked bonds." Bank of England, Financial Stability Paper no. 39.

Blanchard, O., Mauro, P., Acalin, J. (2016) "The case for growth-indexed bonds in advanced economies today." PIIE Policy Brief 16-2.

Borensztein, E., Mauro, P. (2004) "The case for GDP-indexed bonds." Economic Policy, 19(38), 166–216.

Brooke, M., Mendes, R., Pienkowski, A., Santor, E. (2013) "Sovereign default and state-contingent debt." Bank of England, Financial Stability Paper no. 27.

Cabrillac, B., Gauvin, L., Gossé, J.-B. (2017) "Benefits of GDP-indexed bonds for issuing countries, investors and international financial stability". http://voxeu.org/article/benefits-gdp-indexed-bonds.

Carnot, N., Summer, S. (2017) "GDP-linked Bonds: Some Simulations on EU Countries." European Commission Discussion Paper 073, December.

Cecchetti, S., Schoenholtz, K. (2017) "GDP-linked bonds: A primer." http://voxeu.org/article/gdp-linked-bonds-primer

Chamon, M., Mauro, P. (2006) "Pricing growth-indexed bonds". Journal of Banking and Finance, 30(12), 3349–3366.

Cruces, J.J., Trebesch, C. (2013) "Sovereign Defaults: The Price of Haircuts." American Economic Journal: Macroeconomics, 5(5), pp.85–117.

Eichengreen, B. (2003) "Restructuring sovereign debt." Journal of Economic Perspectives, 17(4), 75–98.

Forni, L., Palomba, G., Pereira, J., Richmond, C. (2016) "Sovereign Debt Restructuring and Growth." IMF Working Paper no. 16/147.

G20 (2016) "Communiqué: G20 Finance Ministers and Central Bank Governors Meeting," 23-24 July 2016, Chengdu, China. `www.pbc.gov.cn/english/130721/3106268/index.html`.

G20 (2017) "Compass for GDP-linked bonds." March. `https://goo.gl/9xVnYW`.

Government of Singapore, Ministry of Finance (2008). `www.mof.gov.sg/ers/index.htm`.

International Monetary Fund (2015) "Now is the time: fiscal policies for sustainable growth." IMF Fiscal Monitor, April 2015, 21–49.

International Monetary Fund (2017) "State-contingent debt instruments for sovereigns." IMF Policy Paper, May.

Kamstra, M.J., Shiller, R.J. (2009) "The case for trills: giving the people and their pension funds a stake in the wealth of the nation." Cowles Foundation Discussion Paper no. 1717.

Krugman, P. (1988) "Financing vs. forgiving a debt overhang." Journal of Development Economics, 29, 253–268.

Miyajima, K. (2006) "How to evaluate GDP-linked warrants: price and repayment capacity." IMF Working Paper no. 06/85.

Pirian, A., (2003) "Bulgarian Brady Bonds and the External Debt Swap of Bulgaria." NISPAcee occasional paper in Public Administration and Public Policy, IV(2), pp.3–15.

Sharma, K., Griffith-Jones, S. (2006) "GDP-indexed bonds: making it happen." United Nations, DESA Working Paper no. 21.

Shiller, R. (1993) "Macro markets: creating institutions for managing society's largest economic risks." Oxford, Clarendon Press.

Stumpf, M.H. (2010) "Reflections on the Bosnia Debt Restructuring." Law and Contemporary Problems, 73(4), 301–315.

Trebesch, C., Papaioannou, M., Das, M. (2012) "Sovereign Debt Restructurings 1950-2010: Literature Survey, Data, and Stylized Facts." IMF Working Paper no. 12/2013.

Zettelmeyer, J., Trebesch, C., Gulati, M. (2013) "The Greek debt restructuring: an autopsy." Economic Policy, 28(75), 513–563.

Appendix. Characteristics of bonds with GDP-indexed features issued so far

Issuer country	Main features
Bulgaria[1]	- As a consequence of Bulgaria's (external) debt crisis. - In 1994 Bulgaria signed a Brady contract for the reduction and restructuring of its debt. - Within the restructuring deal there was a clause for recovery of the value and payment was triggered if: (a) current GDP was equal or higher than 125% of GDP in 1993 and (b) there was a GDP increase compared to the previous year. - If those conditions were met, the extra interest rate would be half of the GDP percentage increase (paid in the addition to the underlying plain vanilla coupon). - According to (Miyajima 2006) the source of reference data and GDP measurement units is "ambiguous" and the corresponding term sheet is not clear in the units of measurement. - These warrants were 'callable' and were inseparable from the plain vanilla bonds.
Bosnia and Herzegovina[2]	- In the sequence of the war in Bosnia (1992-1995) that, among other disastrous consequences, led to a significant fall in GDP. The country inherited a legacy of disadvantageous conditions from Yugoslavia, among which, (partially) a considerable high external debt. - In 1997 an agreement on the debt restructuring was achieved and a GDP-performance bond was "settled". - According to (Miyajima 2006) payment on these GDP-warrants would be triggered if: (a) GDP would hit a predetermined target level and would remain at such level for two years and (b) GDP per capita would rise above US$2.80 in 1997 units, adjusted for German consumer price inflation - Also according to the same author, this instrument suffered from poor design and low quality data. - As the Bulgarian GDP-linked warrants (GLWs), were also inseparable from the plain vanilla bonds.
Singapore[3]	- Issuance to low-income citizens of two sets of shares linking payments to GDP-growth (neither tradable nor transferable and could be exchanged only for cash with the government). - The first share – the New Singapore Shares (NSS) – was introduced in 2011 with the purpose of helping the lower-income group during economic downturns. - It consists of annual dividends (on outstanding balances) in the form of bonus shares with a guaranteed 3% minimum rate. An extra dividend, when applicable, corresponds to the real GDP growth rate (if positive) of the previous year.

[1] (Pirian 2003), (Miyajima 2006).
[2] (Stumpf 2010), (Miyajima 2006).
[3] (Government of Singapore - Ministry of Finance 2008), (Miyajima 2006).

	- The second share – the Economic Restructuring Shares (ERS) – was issued with the aim of subsidizing citizens given the Goods and Services Tax increase from 3% to 5%. - Calculation of bonuses is similar to the one of NSS.
Argentina[4]	- Following a period of a severe economic and financial crisis, Argentina defaulted on its sovereign debt obligations by US$82 billion. - After a period of hard negotiations with bondholders of the defaulted debt, in 2005 a debt restructuring was accepted by 76% of them, leading to a bond exchange of US$62 billion in principal. - It included 30-year GLWs that were attached, for a period of 180 days, to the new bonds. - GLWs had no principal and, after the above-mentioned period, could act as "series of standalone, state-contingent coupons". - These instruments were issued in different countries and currencies. - The GLWs would pay annually 5% of excess GDP (defined as the difference between actual real GDP and Base Case real GDP, converted to nominal pesos[5]) if all the following conditions were to be met: (a) actual GDP, expressed in constant peso terms as of the reference date (the year before the one in which payments occur) exceeds the Base Case GDP; (b) the annual growth rates of actual GDP, expressed in constant peso terms as of the reference date, also exceed the Base Case GDP for that year. The growth rate was set at 4.3% for 2005, declining thereafter, reaching 3% from 2015 to 2034; and (c) total cumulative payments should not exceed a payment limit of 48 cents per dollar of notional amount.
Greece[6]	- The Greek sovereign debt crisis led to the 2012 debt restructuring, which included a debt relief of over 50% of that year's GDP. - Within the restructuring package, the new bonds included a set of detachable GDP-linked securities, which could yield an increase in the coupon of up to 1%[7] if (a) nominal GDP in the previous year equals or exceeds the Reference Nominal GDP; (b) real GDP growth equals or exceeds the Reference Real GDP Growth Rate; (c) real GDP growth equals or exceeds 0. - The warrants have a face value, which first equals the face value of the new bond and is reduced by about 5% per year from 2024 to 2042. The principal is used to determine the annual payments, i.e., holders are not entitled to receive it. - The warrants are callable from 2020 on, based on a trailing 30-day market price.

[4] (Benford et al. 2016), (Miyajima 2006).
[5] Excess GDP =(0.05 Excess GDP) x unit of currency coefficient.
[6] (Zettelmeyer et al. 2013).
[7] Payment amount = [1.5 (Real GDP Growth Rate – Reference Real GDP Growth Rate)] x Notional

Term premia dynamics in the US and Euro Area: who is leading whom?

Nikolay Iskrev
Banco de Portugal

January 2018

Abstract
This article examines the dynamic relationship between term premia in euro area and US government bond yields. The term premia are extracted using an affine term structure model using daily data on zero-coupon bond yields. The results show strong co-movement between changes in the premia, especially at the long end of the yield curves. A further investigation of the causal relationship between the euro area and US term premia reveals that only a small fraction of the co-movements can be attributed to one region driving the other. (JEL: G12, E43)

Introduction

While interest rates at all maturities play a role in the borrowing and lending decisions of businesses and households, longer-term rates are typically the ones that matter the most for aggregate spending in the economy. In particular, long-term rates play a central role when businesses decide whether to start new investment projects, households – whether and when to purchase a new home or car, and policy makers – in deciding how to finance government expenditures. From a theoretical point of view, longer-term rates can be seen as risk-adjusted averages of expected future short-term rates. This link between short and long-term rates explains how the transmission mechanism of monetary policy usually works – changes in the short-term interest rate, which is under central banks' direct control, influence aggregate spending decisions by affecting expectations about future short-term rates and thereby changing longer-term rates.[1]

Acknowledgements: I would like to thank Isabel Correia, Nuno Alves, António Antunes, Sandra Gomes, Miguel Gouveia, and the seminar participants at the Bank of Portugal for helpful comments and discussions. The views expressed here are our own and do not necessarily reflect the views of the Bank of Portugal or the Eurosystem.

E-mail: nikolay.iskrev@bportugal.pt

1. In the case of the US Federal Reserve, promoting "moderate long-term interest rates" is one of the explicitly mandated goals, alongside maximum employment and stable prices.

The need to account for risk makes matters more difficult. Both the amount of risk in long-term bonds and its price change over time, giving rise to a time-varying term premium which complicates the relationship between policy rates and long-term rates. The term premium represents the compensation investors in long-term bonds require for the risk that future short rates do not evolve as expected. Given its importance, there has been a large amount of research directed at characterizing the term premium and the factors affecting its level and dynamics.

In this article, I study the relationship between term premia in the yields of euro area (EA) and US government bonds. It is a well-known empirical fact that interest rates of government bonds of advanced economies tend to move closely together, especially at the longer end of the yield curve. One of the objectives here is to establish whether this is also true for the term premium components of the yields. To that end, I estimate affine term structure models of the interest rates for the euro area and the US, and use them to separate expectations from term premia. Then, I measure the degree of co-movement between the levels and the changes in the term premia using linear correlation coefficients. The second objective of the article is to explore the evidence for a causal relationship between the two term premia, that is, the extent to which we can say that movements in the term premia of one economic area drive the movements in the term premia of the other area. For that purpose I estimate static and dynamic versions of indicators that have been proposed in the time series literature to measure the strength and direction of causal relationships. The results from this analysis show that there exist time-varying causal linkages between the EA and US term premia. At the same time, it is found that only a relatively small fraction of the observed co-movements can be attributed to one region driving the other.

The rest of the article is organized in four sections. The first one presents some basic yield curve concepts and introduces the expectations theory of interest rates. The second section first outlines and estimates an affine term structure model, which is used to decompose long-term yields into expectations and term premia, and then evaluates the strength of co-movement between euro area and US term premia. The third section describes and estimates several measures of causality between the term premia. The last section offers some concluding remarks.

Term structure of interest rates

This section introduces some basic yield curve terminology and presents the expectations theory of interest rates, which is in the background of most modern term structure models.

Notation and basic concepts

While bonds typically pay coupons during their lifetime, economists prefer to work with zero-coupon bonds, also known as pure discount bonds. These are bonds that promise to pay one euro on a given future day – the maturity date of that bond. Non-zero coupon bonds can be seen as portfolios of zero-coupon bonds. The interest rates on the zero-coupon bonds are called yields, and the function describing the relationship between bond maturities and their yields at a given point in time is called the yield curve. Zero-coupon bonds are convenient because there exists a simple relationship between the price $P_t^{(n)}$ at time t and the yield $y_t^{(n)}$ at time t of such bonds:

$$P_t^{(n)} = e^{-n \times y_t^{(n)}},$$

where n is the time to maturity measured in years. The yield is the continuously compounded annualized return from holding the zero coupon bond until maturity. At a given point in time the yield of a bond will depend on its maturity, and the yield curve is the function describing that relationship. Figure 1 shows several historical yield curves for maturities between 3 months and 10 years for the euro area and the US. The observations are from the first and last months in our sample – from October 2004 until October 2017. Also shown are the average curves over the sample period. Several features of the figure are worth noting: first, the curves are upward sloping and have very similar shapes, both across time and regions. Upward-sloping yield curves are more common in general although historically there have been episodes of downward-sloping curves, for instance the US in the early 2000s. Second, both the EA and US yield curves have shifted downwards over the sample period, and remain below the average curves at the end of the sample. However, while at the beginning of the sample period the levels of the EA and US yield curves are approximately the same, they are very different at the end of the sample, with the EA yield curve being much lower than the one for the US. Explaining such differences in the shape of the yield curve across time and economic regions is one of the main objectives of the research on the term structure of interest rates.

The expectations hypothesis

The expectations theory of interest rates is among the oldest and most popular models of the term structure.[2] In its general form, the expectation hypothesis postulates that long-term rates and expected short-term rates must be linked.[3]

2. The main ideas behind the expectations hypothesis can be traced back to the work of Fisher (1896) and Lutz (1940).

3. In the literature it is common to distinguish between the "pure expectations hypothesis", which states that the long rates are equal to the average expected short rates, and the

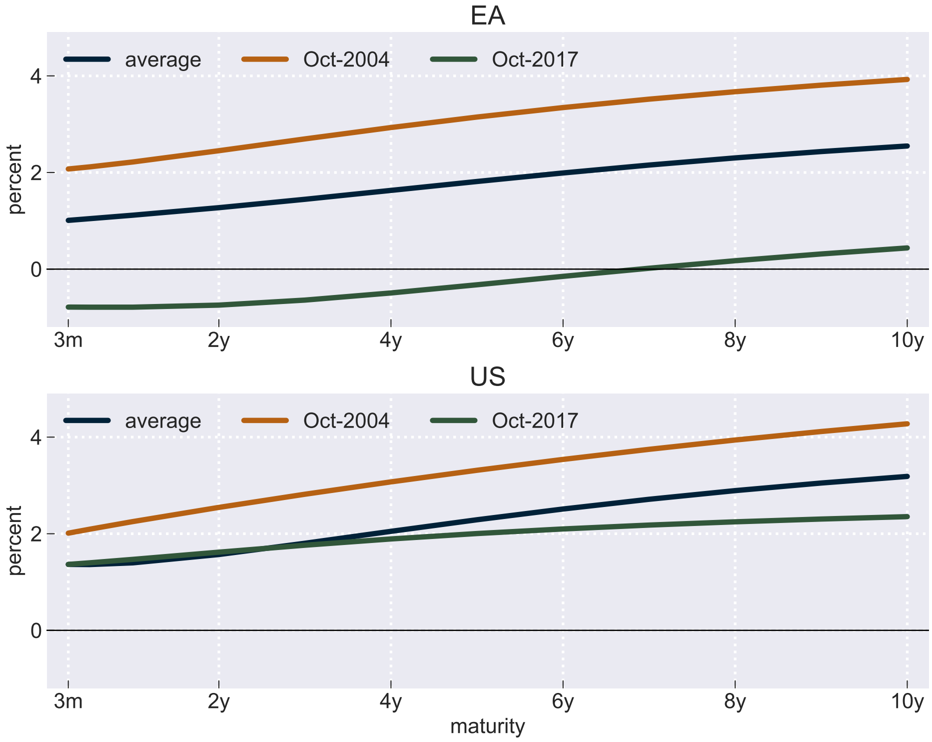

FIGURE 1: **EA and US yield curves**. The figure shows the EA and US zero-coupon yield curves at the beginning and the end of the sample (October 2004 and October 2017, respectively), as well as the average yield curves across the sample period. Source: ECB, FRB, and own calculations.

The theory is motivated by the observation that investors choose between short and long-term bonds by comparing the return of the long-term bond to the expected return of an investment strategy of rolling-over a sequence of short-term bonds. To understand the basic intuition, assume for a moment that future yields are certain, and consider an investor who chooses between two investment strategies: buying 2-year bonds today, or buying 1-year bonds today, the proceeds from which are then re-invested in 1-year bonds one year hence. Using the first strategy, the investor has to pay $P_t^{(2)} = e^{-2\times y_t^{(2)}}$ euros today to receive one euro in two years. The price next year of a 1-year bond is $P_{t+1}^{(1)} = e^{-y_{t+1}^{(1)}}$. The price today of $P_{t+1}^{(1)}$ one-year bonds is $P_t^{(1)} \times P_{t+1}^{(1)} = e^{-y_t^{(1)}} P_{t+1}^{(1)}$. Therefore, to receive 1 euro in two years using the second strategy, the investor has to pay $e^{-y_t^{(1)}} \times e^{-y_{t+1}^{(1)}}$ today. The two strategies yield the same

"expectations hypothesis" which states that deviations of long rates from the average expected short rates are constant over time.

return and therefore must require the same initial investment, i.e

$$e^{-2\times y_t^{(2)}} = e^{-(y_t^{(1)}+y_{t+1}^{(1)})}$$

Hence, absence of arbitrage requires that

$$y_t^{(2)} = \frac{1}{2}(y_t^{(1)} + y_{t+1}^{(1)})$$

Using the same argument, we can establish the following relationship between the yield of bonds with n years to maturity and the yield on the present and future one-year bonds:

$$y_t^{(n)} = \frac{1}{n}\left(y_t^{(1)} + y_{t+1}^{(1)} + \ldots + y_{t+n-1}^{(1)}\right) \tag{1}$$

Uncertainty about future short-term yields means that investment decisions have to be made on the basis of investors' expectations about future yields. Furthermore, investors are averse to risk and will demand a premium for holding riskier long-term bonds. The classical formulations of the expectations hypothesis set the premium to zero or to a non-zero constant. However, numerous studies testing formulations of the expectations hypothesis have found evidence for time-varying risk premia (see for instance Mankiw *et al.* (1984), Fama and Bliss (1987), Campbell and Shiller (1991)). This leads to the following more general representation of bond yields:

$$y_t^{(n)} = \frac{1}{n}\sum_{h=0}^{n-1} E_t y_{t+h}^{(1)} + TP_t^{(n)}, \tag{2}$$

where $TP_t^{(n)}$ denotes the term premium at time t for bonds with n years to maturity. In order to separate the term premia from the expectations component, we need a model for the term structure. The next section describes and estimates one such model.

Yield decomposition based on affine term structure model

In this section, I use daily zero-coupon yields data to decompose observed long-term rates into expectation components and term premia. To that end, I estimate a no-arbitrage affine term structure model of the interest rates. According to this model, both the actual yields and the expectation components can be expressed as affine functions of a small number of risk factors, which are modeled as linear processes. Ruling out arbitrage opportunities imposes restrictions on the yields' behavior over time and across different maturities. Those restrictions facilitate the estimation of the model in terms of a small number of parameters. A fuller description of the affine term structure model and its derivation are presented in the Appendix.

Data and estimation

I estimate the affine term structure model using daily zero-coupon yields for the EA and the US. To compute the daily yield curves I use the Svensson (1994) model with parameter estimates provided by the ECB and the US Federal Reserve.[4] In the case of the EA the yields are of AAA-rated sovereign bonds, which are comparable in terms of risk properties to the US treasury bonds.[5] Using the estimated parameters I construct daily yield curves for maturities from 1 month up to 10 years, for the period between September 2004 and October 2017.[6] The time series of the EA and US zero-coupon yields for selected maturities are presented in Figure 2.

I estimate the model outlined above following a procedure developed by Adrian *et al.* (2013) (ACM henceforth), who show that the underlying model parameters can be estimated using a series of linear regressions. Specifically, their approach takes the risk factors to correspond to the first few principal components of the observed bond yields, and models the factors as a standard vector autoregressive model. The parameters of the model are then obtained in three steps using standard OLS regressions. More details on the estimation procedure is provided in the Appendix.

Number of risk factors

Following the work of Litterman and Scheinkman (1991), it is common in the literature to summarize the term structure using principal components of the covariance matrix of the zero-coupon yields. Typically, it is found that the first three principal components are sufficient to capture most of the variation in the yields. In other words, there are three significant risk factors explaining the shape of the yield curve. These factors are typically referred to as *level*, *slope* and *curvature* factors. The reason for these labels can be understood by considering the factor loadings displayed in Figure 3. The factor loadings show how sensitive yields at different maturities are to changes in each principal component, or risk factor. In the figure we see that changes in the first factor result in a level shift for the yields of all maturities. Changes in the second factor move the short and long maturities in opposite directions.

4. The estimated parameters are downloaded from `http://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/euro_area_yield_curves/html/index.en.html` for the EA and `https://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html` for the US. The Svensson model is also used by the ECB to produce daily yield curves for the EA, as well as by Gürkaynak *et al.* (2007) whose zero-coupon yield data set is commonly used for estimating term structure models for the US.

5. Note that the selection of EA countries with AAA rating changes over time. The ratings ECB uses are provided by Fitch Rating.

6. Official data for the EA is available from 6 September 2004, while the data for the US starts in 14 June 1961.

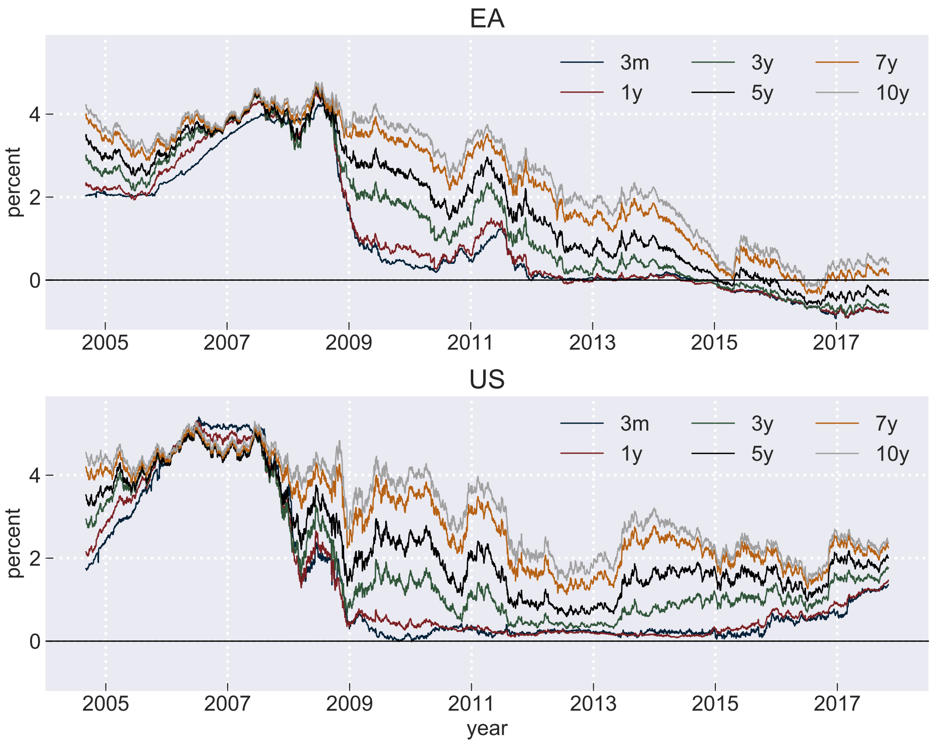

FIGURE 2: **EA and US zero-coupon yields**. The figure shows the time series of EA and US zero-coupon yields for selected maturities.
Source: ECB and FRB.

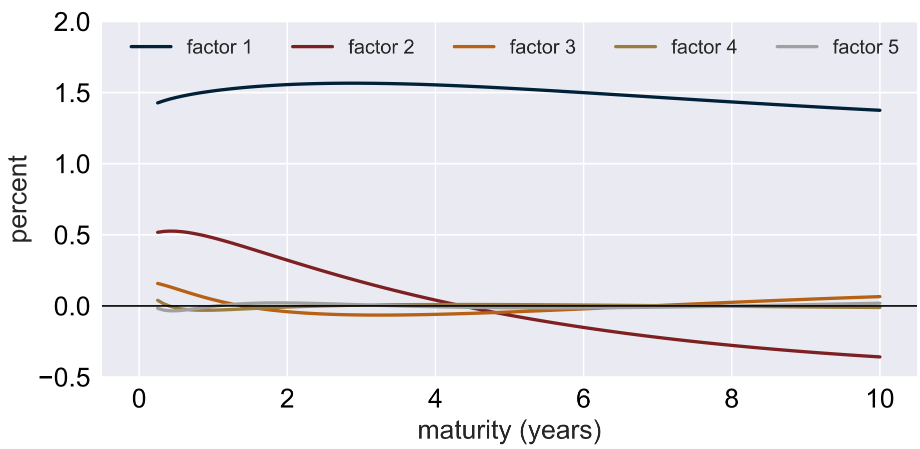

FIGURE 3: **Risk factors loadings**. The figure displays the loadings of bond yields on the first five principal components.
Source: ECB, FRB, and own calculations.

	all	3m	1y	2y	3y	4y	5y	10y
# of PCs				(a) EA				
1	96.439	86.158	89.698	95.056	98.247	99.650	99.855	94.282
2	3.434	12.473	10.147	4.884	1.610	0.208	0.054	5.529
3	0.115	1.184	0.141	0.038	0.139	0.138	0.081	0.168
4	0.009	0.069	0.004	0.021	0.002	0.003	0.009	0.019
5	0.003	0.105	0.011	0.000	0.002	0.001	0.000	0.002
# of PCs				(b) US				
1	94.685	85.601	89.251	92.898	96.084	98.451	99.573	88.355
2	4.972	11.552	10.101	7.033	3.725	1.164	0.030	11.053
3	0.309	2.296	0.643	0.004	0.165	0.383	0.383	0.537
4	0.032	0.525	0.001	0.064	0.025	0.000	0.014	0.051
5	0.002	0.023	0.004	0.001	0.001	0.002	0.000	0.004

TABLE 1. Percent of the variance explained by the first 5 principal components.

Source: Own calculations.

Lastly, changes in the third factor move the short and long maturities in the same direction, leaving the medium-term maturities mostly unaffected. In addition, the figure shows that the yields of all maturities are mostly sensitive only to the first three factors, while changes in either the fourth or the fifth principal component have only a minor impact. Figure 3 is based on data for the EA, but the results with US data are very similar.

Another standard approach for determining the number of factors is to compute the fraction of the total variance of the observed yields explained by each additional risk factor. As can be seen in Table 1, for both the EA and the US, the first three principal components are sufficient to capture more that 99% of the variance of the yields as a whole, as well as the variances of yields at selected maturities.

These results are in line with the broad consensus in the literature that the first three principal components of the yield curve are sufficient to capture well the dynamics of the term structure. However, the ACM estimates of the US term premia are based on five pricing factors, and that is the specification underlying the yield curve decomposition published by the New York Fed. For consistency with their approach, here I present results based on a five factor model for both the EA and US yield curves.[7]

7. It should be noted that the US term premia estimates published daily by the New York Fed are estimated with a sample starting in 1961, while the estimates presented in this article are obtained with a sample starting in 2004. The main impact this difference has on the results is on the level of term premium, which is higher with the more recent sample. The dynamics of the term premia remains almost unchanged. This level effect is due to the fact that the mean of the short-term rate is much higher in the longer sample, which drives the expectations component up and the term premium down.

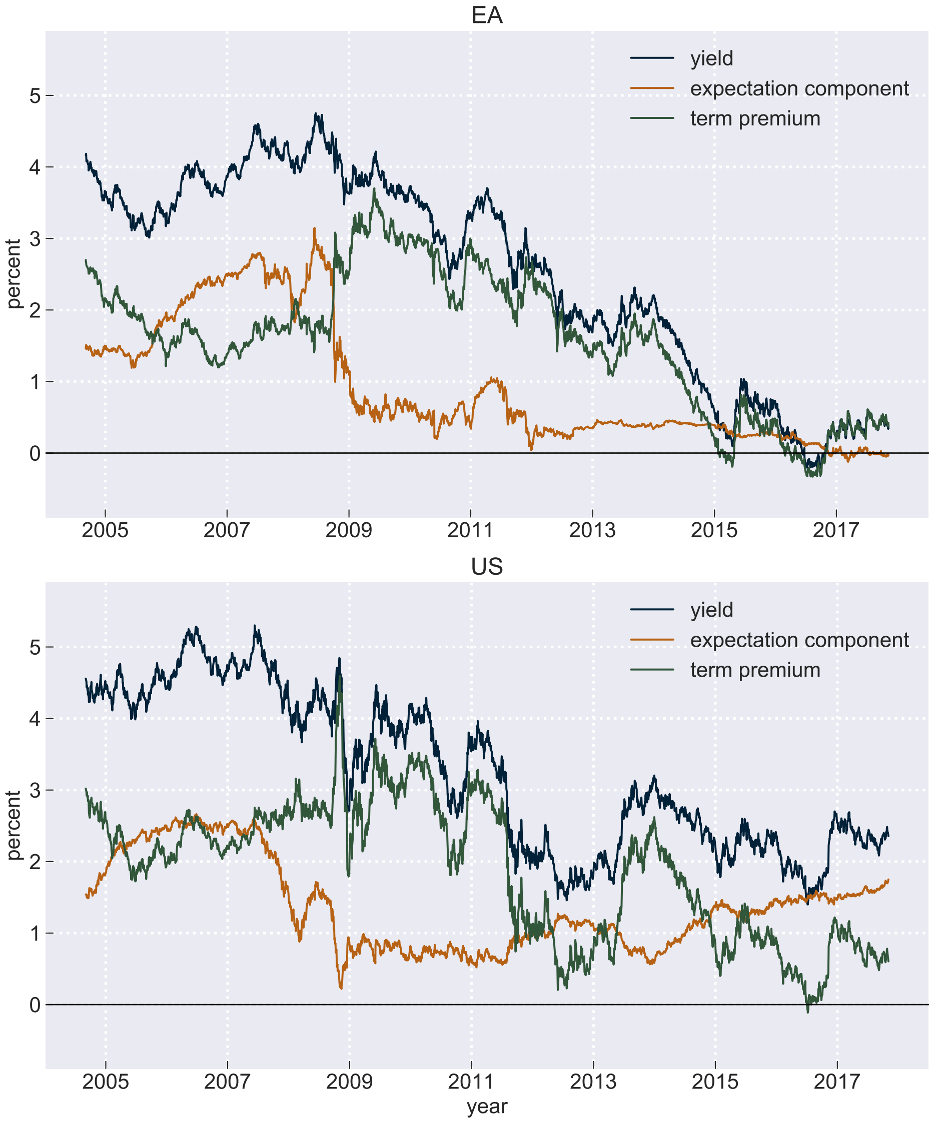

FIGURE 4: **10 year yield decomposition**. This figure plots decompositions of the EA and US 10-year daily yields into expectation components and term premia.
Source: ECB, FRB, and own calculations.

Term premia estimates

Following ACM, I estimate the parameters of the model using end-of-month observations of the zero-coupon yields. Given the estimated parameters, I can compute the model-implied decomposition of the fitted yields $y_t^{(n)}$ into expectations component $\tilde{y}_t^{(n)}$ and term premium $TP_t^{(n)}$ for all maturities and at any point in time. In particular, with daily observations of the risk factors, extracted as principal components of the daily zero-coupon yields, I can decompose the yields into expectations component and term premia at daily frequency. Figure 4 shows an example with daily decompositions of the 10-year bond yields in the EA and US. In the case of the EA yields, for instance, the decomposition suggests that the return of the 10-year yields into positive territory at the end of 2016 was entirely due to an increase in the term premium, i.e. the compensation for holding longer-term bonds by investors. In fact, the 10-year yields have tracked closely the movements in the term premium for most of the time since 2012, due to the expectation component remaining relatively flat over that period. On the other hand, the expectations component in the US 10-year yields has been increasing steadily since 2014. This rise in the short rate expectations explains to a large extent the observed divergence in the 10-year yields in the two regions. At the same time, as can be seen better in Figure 5, the 10-year term premia in the EA and the US have followed very similar paths during the sample period. In both regions the term premia reached historically low levels in the second half of 2016. Also shown in the figure is the 250-day rolling correlation between the two series. During most of the period the correlation is positive and very strong, often in excess of 0.9.

However, using correlation here may be misleading since the two series appear to be non-stationary.[8] Thus, it is more reasonable to compare changes in the term premia components of the respective bond returns. Figure 6 shows the changes in the 10-year term premia in the EA and the US and the 250-day rolling correlation between those series. Again, during most of the sample period the correlation is positive and relatively strong. This is not a feature of the 10-year term premia only. Figure 7 shows a heat plot of rolling correlations between changes in the EA and US term premia for all maturities up to 10 years. The degree of correlation tends to be stronger for longer maturities, and is about as high as for the 10-year premia for all maturities above 6 or 7 years. On the other hand, for maturities of less than 4 years the correlation tends to be week and is sometimes even negative.

8. This observation is confirmed by formal unit root tests the results of which are presented in the Appendix.

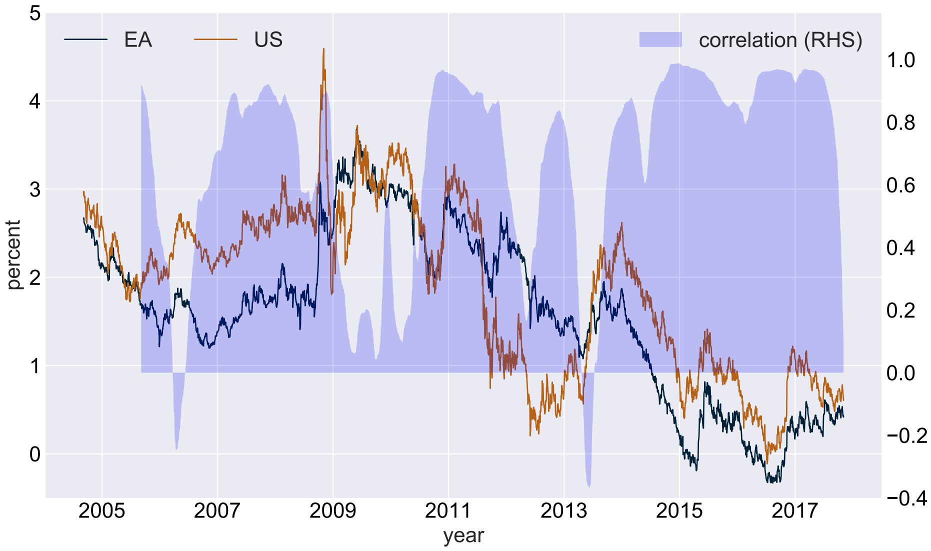

FIGURE 5: **10-year EA and US term premia**. The figure shows 10-year EA and US term premia and 250-day rolling pairwise correlations between the two series.
Source: ECB, FRB, and own calculations.

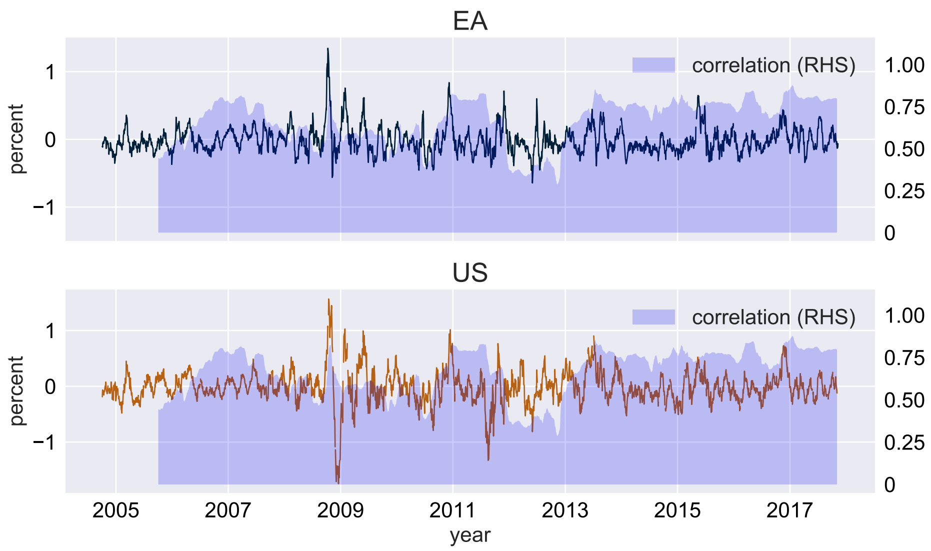

FIGURE 6: **Changes in the 10-year EA and US term premia**. The figure shows the changes in the 10-year EA and US term premia and 250-day rolling pairwise correlations between the two series.
Source: ECB, FRB, and own calculations.

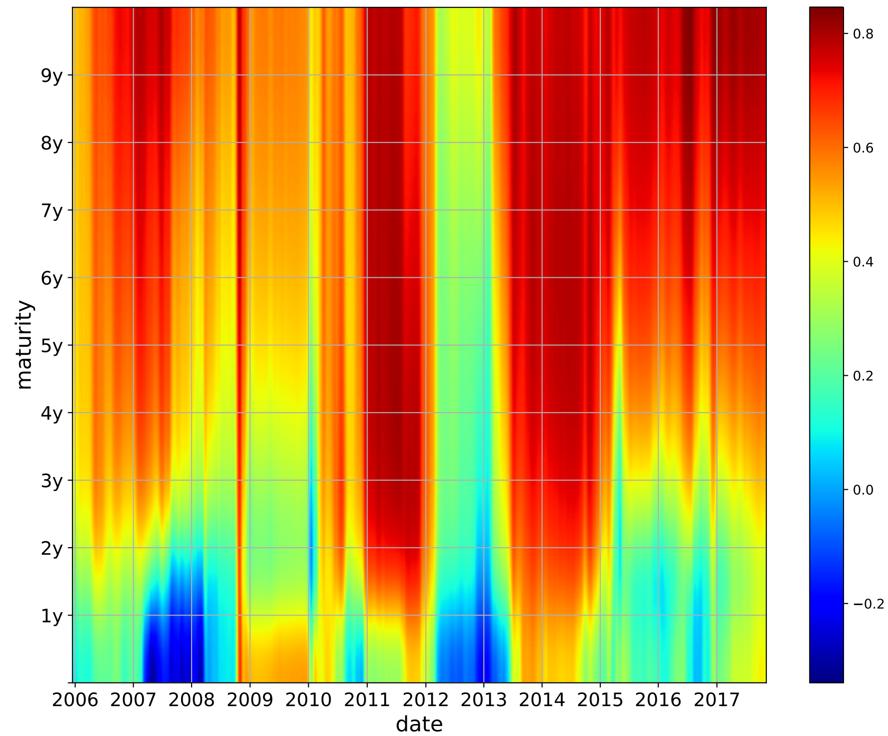

FIGURE 7: **Rolling correlations between changes in the EA and US term premia**. The figure show 250-day rolling pairwise correlations between changes in the EA and US term premia for all maturities up to 10 years.
Source: ECB, FRB, and own calculations.

Detecting and measuring directionality

Indicators

The results in the previous section show that changes in the term premia in the EA and US are strongly positively correlated, especially at the longer end of the yield curve. In this section I consider the evidence for directionality in the interactions between the two variables. Specifically, I estimate three indicators designed to detect and quantify the strength of causal interaction in time series. The indicators are Granger causality, transfer entropy and directional connectedness, and are described below.

Granger causality. Stated simply, the definition of Granger causality is that a variable X *causes* a variable Y if a forecast of Y using X is more accurate than a forecast of Y without using X. To make this definition operational, one needs to specify a forecasting model for Y and typically this is done using linear vector autoregressions (VAR). Then, testing for causality amounts to comparing the size of the forecast errors of Y from a VAR which includes lags of X to the size of the errors from a VAR without those lags.

Transfer entropy. The concept of Granger causality can be interpreted in terms of information content, i.e. the past of variable X containing information about the future of variable Y, information not contained in the past of Y itself. From this perspective, one can define a more flexible, i.e. non-linear, model for predicting Y, as well as use a more general measure of information than the reduction of forecast error variance, which underlies the standard approach to testing for Granger causality. This is in essence what the concept of transfer entropy tries to accomplish.[9] The amount of information from X to Y is measured as the reduction of uncertainty about the future of Y using a model-free measure, namely the entropy of the empirical distribution of the data.

Directional connectedness. In a series of papers, Diebold and Yilmaz (2009, 2012, 2015) developed a measure of connectedness for the purpose of assessing the strength and direction of interdependence across financial markets in different countries. The measure is based on variance decompositions estimated from VAR applied to two or more financial variables. In particular, the connectedness from X to Y is determined by the

9. The entropy of a variable is defined as the negative expected value of the logarithm of the probability distribution of that variable. In the case of a normally distributed variable, the entropy is equivalent to the variance of that distribution. Transfer entropy, as a measure of the amount of information transferred from one time series process to another, was introduced by Schreiber (2000)

share of the forecast error variance of Y due to shocks in X. The identification of the shocks is achieved using the generalized variance decomposition approach of Pesaran and Shin (1998).

Similar to the Granger causality measure, the notion of connectedness can be interpreted in terms of information content, namely, the amount of additional information about future values of one variable contained in the shocks associated with another variable. As before, information is quantified as the reduction of uncertainty about the future values of the first variable. Instead of information in the second variable itself, connectedness is about the impact of the shocks associated with that variable. This common interpretation suggests that we can use the following general representation of the three measures:

$$I_{X \to Y} = 100 \times \left(1 - \frac{Uncertainty(Y|X,Z)}{Uncertainty(Y|Z)}\right) \tag{3}$$

Note that having more information cannot increase uncertainty. Therefore, $Uncertainty(Y|X,Z) \leq Uncertainty(Y|Z)$ is always true. Equality would imply that X contributes no information about Y, once Z is observed. In that case $I_{X \to Y} = 0$. On the other extreme, we could have $Uncertainty(Y|Z) > Uncertainty(Y|X,Z) = 0$, which means that observing both X and Z is equivalent to also observing Y. In that case we have $I_{X \to Y} = 100$.

In the case of both Granger causality and transfer entropy, Y represents future values of one observed variable, for example the 10-year EA term premium, X represents the past values of the other observed variable, i.e. the 10-year US term premium, and Z represents the past values of the first observed variable – the 10-year EA term premium. The value of the indicator in both cases shows the reduction of uncertainty about the future values of the 10-year EA term premium as a result of observing the past values of the 10-year US term premium, compared to using only the past values of the 10-year EA term premium. The difference between these two indicators is in how uncertainty is estimated – with a VAR model and using the forecast error variance in the case of Granger causality, and with a non-parametric estimator of entropy – in the case of transfer entropy. For the directed connectedness measure, Y is again the future values of an observed variable – the 10-year EA term premium – but X represents the future values of the shock associated with the other variable, i.e. the 10-year US term premium, while Z represents the past values of both observed variables, EA and US 10-year term premia.

Results

I estimate the measures of directionality using both the full sample and rolling-window samples. The full sample results are presented in Table 2. Two of the measures – the Granger causality and the directional connectedness – indicate a stronger causal impact from the US to the EA term premia changes. The transfer entropy shows the inverse relationship, i.e. the EA having stronger impact. All three measures agree that the causal influence from one area to the other is relatively weak.

	$EA \rightarrow US$	$US \rightarrow EA$
Granger causality	1.6	2.9
Transfer entropy	4.4	3.6
Directional connectedness	4.4	9.0

TABLE 2. **Static indicators of directional influence**. The values represent the per cent reduction in uncertainty regarding future yields in one area, due to the information from the past yields (in the case of Granger causality and transfer entropy) or future shocks (in the case of directional connectedness) from the other area.

The sample is from September 7, 2004 through October 31, 2017
Source: Own calculations.

To see how the degree of causation changes over time, I perform a rolling-window analysis using windows with a length of 250 days. The results are displayed in Figure 8. They show that the strength of causal influence changes over time, and in some periods the impact from the EA is stronger, while in others the influence from the US dominates. In particular, all three measures are consistent in suggesting that EA has a stronger impact on the US during the period from 2011 through 2013, while from the middle of 2013 until the second half of 2014 the degree of causality from US to EA is stronger. The Granger causality and directional connectedness measures also indicate that influence from the US dominates that from the EA in the beginning of the sample – from 2006 until 2008. In the case of transfer entropy, the EA has somewhat stronger impact during that period.

Overall, with a few exceptions, the transfer entropy measure suggests a relatively more equal degree of causal influence from either area, while the other two measure show several periods where causal influence from one of the areas clearly dominates. At the same time, all three measures indicate a relatively small causal impact from either area to the other. In terms of information transfer, this means that there is a relatively small amount of unique information in either series that helps predict the future developments in the other. Therefore, one of the main reasons for the strong co-movement between the series must be that they are both subject to influence by a global factor or factors. For instance, international factors driving uncertainty about future inflation will also affect term premia in different markets. Empirical evidence linking the downward slope in international term premia to declining inflation uncertainty are discussed by Wright (2011).

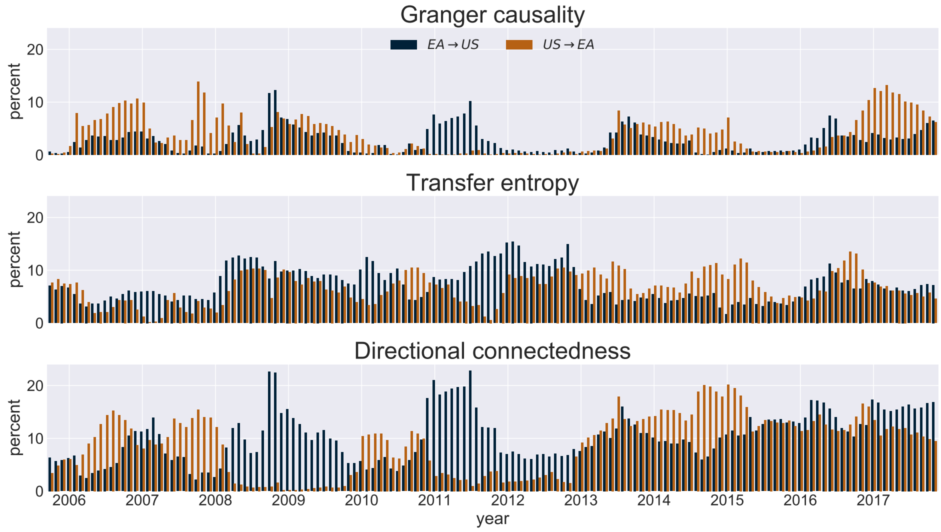

FIGURE 8: **Dynamic indicators of directional influence**. The figure shows 250-day rolling window estimates of the indicators. The values represent the per cent reduction in uncertainty regarding future yields in one area, due to the information from the past yields (in the case of Granger causality and transfer entropy) or future shocks (in the case of directional connectedness) from the other area. Source: ECB, FRB, and own calculations.

Concluding remarks

This article investigated the dynamics of term premia in EA and US government bonds. I found that there is a strong co-movement between the premia, especially at the long end of the yield curve, both in terms of the levels as well the changes in the two series. Further analysis of the potential causal relationship between the bond term premia revealed that only a small fraction of the joint dynamics can be attributed to one region driving the other. This part of the analysis was based on several different indicators which, in contrast to measures of co-movement like correlation, are non-symmetric and provide information about the direction of causality. While all indicators suggest the existence of a time-varying causal linkages between EA and US term premia, they were found to be relatively weak. Given this evidence, a more plausible explanation of the strong co-movement is that there exist a common global factor that affects term premia in both regions.

References

Adrian, Tobias, Richard K. Crump, and Emanuel Moench (2013). "Pricing the term structure with linear regressions." *Journal of Financial Economics*, 110(1), 110–138.

Campbell, John Y and Robert J Shiller (1991). "Yield spreads and interest rate movements: A bird's eye view." *The Review of Economic Studies*, 58(3), 495–514.

Diebold, Francis X and Kamil Yilmaz (2009). "Measuring financial asset return and volatility spillovers, with application to global equity markets." *The Economic Journal*, 119(534), 158–171.

Diebold, Francis X and Kamil Yilmaz (2012). "Better to give than to receive: Predictive directional measurement of volatility spillovers." *International Journal of Forecasting*, 28(1), 57–66.

Diebold, Francis X and Kamil Yilmaz (2015). *Financial and Macroeconomic Connectedness: A Network Approach to Measurement and Monitoring*. Oxford University Press, USA.

Fama, Eugene F and Robert R Bliss (1987). "The information in long-maturity forward rates." *The American Economic Review*, pp. 680–692.

Fisher, Irving (1896). *Appreciation and Interest: A Study of the Influence of Monetary Appreciation and Depreciation on the Rate of Interest with Applications to the Bimetallic Controversy and the Theory of Interest*, vol. 11. American Economic Association.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright (2007). "The US Treasury yield curve: 1961 to the present." *Journal of Monetary Economics*, 54(8), 2291–2304.

Litterman, Robert B and Jose Scheinkman (1991). "Common factors affecting bond returns." *The Journal of Fixed Income*, 1(1), 54–61.

Lutz, Friedrich A (1940). "The structure of interest rates." *The Quarterly Journal of Economics*, 55(1), 36–63.

Mankiw, N Gregory, Lawrence H Summers, *et al.* (1984). "Do Long-Term Interest Rates Overreact to Short-Term Interest Rates?" *Brookings Papers on Economic Activity*, 15(1), 223–248.

Pesaran, H Hashem and Yongcheol Shin (1998). "Generalized impulse response analysis in linear multivariate models." *Economics letters*, 58(1), 17–29.

Schreiber, Thomas (2000). "Measuring information transfer." *Physical review letters*, 85(2), 461.

Svensson, Lars EO (1994). "Estimating and interpreting forward interest rates: Sweden 1992-1994." Tech. rep., National Bureau of Economic Research.

Wright, Jonathan H (2011). "Term premia and inflation uncertainty: Empirical evidence from an international panel dataset." *The American Economic Review*, 101(4), 1514–1534.

Appendix: Arbitrage-free Gaussian affine term structure models

Affine term structure models model zero-coupon bond yields as functions of a vector of variables X_t, called pricing or risk factors, and assumed to follow a Gaussian vector autoregression (VAR(1)):

$$X_t = \mu + \Phi X_{t-1} + \varepsilon_t, \;\; v_t \sim N(0, \Sigma) \tag{A.1}$$

Let $P_t^{(n)}$ be the price of a zero-coupon bond with maturity n at time t. Assuming that there is no arbitrage implies the existence of a price kernel M_t such that

$$M_t = E_t\left(M_{t+1} P_{t+1}^{(n-1)}\right) \tag{A.2}$$

Assume that the pricing kernel is exponentially affine, i.e:

$$M_t = \exp\left(-r_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\Sigma^{-1/2}v_{t+1}\right) \tag{A.3}$$

where $r_t = -\ln(P_t^{(1)})$ is the continuously compounded one-period rate, and λ_t are the market prices of risk. Both r_t and λ_t are assumed to be affine functions of the pricing factors

$$r_t = \delta_0 + \delta_1 X_t \tag{A.4}$$

$$\lambda_t = \Sigma^{-1}\left(\lambda_0 + \lambda_1 X_t\right) \tag{A.5}$$

Denote with $rx_{t+1}^{(n-1)}$ the log of the excess holding return of a bond maturing in n periods:

$$rx_{t+1}^{(n-1)} = \ln P_{t+1}^{(n-1)} - \ln P_t^{(n)} - r_t \tag{A.6}$$

ACM show that if $\{rx_{t+1}, v_{t+1}\}$ are jointly normally distributed, then

$$E_t\left(rx_{t+1}^{(n-1)}\right) = \beta^{(n-1)}\left(\lambda_0 + \lambda_t X_t\right) - \frac{1}{2}var\left(rx_{t+1}^{(n-1)}\right) \tag{A.7}$$

where $\beta^{(n-1)} = cov\left(rx_{t+1}^{(n-1)}, v_{t+1}'\right)\Sigma^{-1}$. Furthermore, the return generating process for the log excess returns is

$$\begin{aligned} rx_{t+1}^{(n-1)} &= \beta^{(n-1)}\left(\lambda_0 + \lambda_t X_t\right) - \frac{1}{2}\left(\beta^{(n-1)\prime}\Sigma\beta^{(n-1)} + \sigma^2\right) \\ &+ \beta^{(n-1)\prime}v_{t+1} + e_{t+1}^{(n-1)} \end{aligned} \tag{A.8}$$

where $e_{t+1}^{(n-1)}$ is a return pricing error assumed to follow an i.i.d. process with mean 0 and variance σ^2. The above equation can be written in a stacked form for all t and n as follows

$$\begin{aligned}\mathbf{rx} &= \boldsymbol{\beta}\left(\lambda_0\iota_T + \lambda_t X_{\_}\right) - \frac{1}{2}\left(\mathbf{B}^* vec(\Sigma) + \sigma^2\iota_N\right)\iota_T' \\ &+ \boldsymbol{\beta}' V + E \end{aligned} \tag{A.9}$$

where $\mathbf{rx}$ is $N \times T$ matrix of excess returns, $\boldsymbol{\beta}$ is $K \times N$ matrix of factor loadings, ι_T and ι_N are T and N dimensional vectors of ones, $X_{\_} = [X_0, X_1, \ldots, X_{T-1}]$ is a $K \times T$ matrix of pricing factors, $\mathbf{B}^* = [vec(\beta^{(1)}\beta^{(1)\prime}), \ldots, vec(\beta^{(N)}\beta^{(N)\prime})]$ is an $N \times K^2$ matrix, V is a $K \times T$ matrix, and E is an $N \times T$ matrix.

A.1. Estimation

ACM show that the parameters of the model can be obtained using a series of linear regressions. We start by estimating equation (A.1) by OLS. The estimated innovations $\hat{v}_t$ are stacked into a matrix $\hat{V}$ which is used as a regressor in the estimation of the reduced-form of (A.9) by OLS:

$$\mathbf{rx} = \mathbf{a}\iota_T' + \mathbf{c}X_{\_} + \boldsymbol{\beta}' V + E \tag{A.10}$$

Using the restrictions equation (A.9) imposes on $\mathbf{a}$ and $\mathbf{c}$ in the equation above gives us the following estimates of the risk parameters λ_0 and λ_1:

$$\hat{\lambda}_0 = (\hat{\boldsymbol{\beta}}\hat{\boldsymbol{\beta}}')^{-1}\hat{\boldsymbol{\beta}}\left(\hat{\boldsymbol{a}} + \frac{1}{2}(\mathbf{B}^* vec(\hat{\Sigma}) + \hat{\sigma}^2\iota_N)\right) \tag{A.11}$$

$$\hat{\lambda}_1 = (\hat{\boldsymbol{\beta}}\hat{\boldsymbol{\beta}}')^{-1}\hat{\boldsymbol{\beta}}\hat{\boldsymbol{c}} \tag{A.12}$$

where $\hat{\sigma}^2$ is computed using the estimated residuals of (A.10). Lastly, we estimate the short rate parameters δ_0 and δ_1 by OLS regression of equation (A.4).

A.2. Term premium

The affine structure of the model implies that the continuously compounded yield on a $n-$period zero-coupon bond at time t, defined as $y_t^{(n)} = -\frac{1}{n}\log P_{t,n}$ is given by

$$y_t^{(n)} = -\frac{1}{n}\left(A_n + B_n' X_t\right) \tag{A.13}$$

where the A_n and B_n parameters are derived recursively using the following system of equations:

$$A_n = A_{n-1} + B_{n-1}'(\mu - \lambda_0) + \frac{1}{2}\left(B_{n-1}'\Sigma B_{n-1}' + \sigma^2\right) - \delta_0 \tag{A.14}$$

$$B_n' = B_{n-1}'(\Phi - \lambda_1) - \delta_1' \tag{A.15}$$

$$A_0 = 0, \;\; B_0 = \mathbf{0} \tag{A.16}$$

The yield in (A.13) includes a compensation for risk, demanded by risk-averse investors to invest in a longer-term bond instead of rolling over a series of short-term bonds. That is, we can decompose the model-implied yields into an expectation component and a term premium:

$$y_t^{(n)} = \frac{1}{n}\sum_{j=0}^{n-1} E_t r_{t+j} + TP_t^{(n)} \tag{A.17}$$

where the first term represents the risk-neutral yield, defined as the yield that would be demanded by investors which are risk-neutral. To obtain the risk-neutral yield we set the price-of-risk parameters λ_0 and λ_1 to zero, and use the recursions in (A.14) and (A.15) to derive the risk-adjusted parameters $\tilde{A}_n$ and $\tilde{B}_n$. The risk-neutral yields are computed using:

$$\tilde{y}_t^{(n)} = -\frac{1}{n}\left(\tilde{A}_n + \tilde{B}_n' X_t\right) \tag{A.18}$$

The term premium is obtained as the difference between actual (model-implied) and risk-neutral yield

$$TP_t^{(n)} = y_t^{(n)} - \tilde{y}_t^{(n)} \tag{A.19}$$

A.3. Unit root tests

	EA		US	
	level	**diff.**	**level**	**diff.**
Dickey-Fuller GLS test	-0.18 (-1.95)	-6.83 (-1.95)	-0.77 (-1.95)	-6.97 (-1.95)
Phillips-Perron test	-1.75 (-3.41)	-9.04 (-3.41)	-3.03 (-3.41)	-9.44 (-3.41)

TABLE A.1. **Testing for unit root in the level and differences of the EA and US 10-year term premium.**

The null hypothesis for both tests is that the process contains a unit root. The table shows the values of the test statistics and the respective 5% critical values (in parenthesis).
Source: Own calculations.

www.bportugal.pt