The cyclicality of the Portuguese labour market: a macroeconomic perspective in the OECD context[1]

Pedro Amaral[2]

ABSTRACT

The portuguese labour market's cyclical fluctuations show little correlation with the aggregate business cycle as given by fluctuations in GDP per worker. Even though there are other OECD countries whose labour markets exhibit an equally tenuous relation with the business cycle as Portugal, the norm is a higher correlation. On the other hand, the Portuguese business cycle shows a degree of persistency, or temporal correlation, that ranks among the lowest in the OECD.This article argues that such facts have important implications for macroeconomics models of the labour market.

Introduction

The consequences of the last recession for labour markets in advanced economies around the world were astounding: in Portugal, the unemployment rate hit an unprecedented seventeen percent; in Spain it was almost ten percentage points higher than in Portugal; and even in countries traditionally thought of as having very flexible labour markets, like the U.S., unemployment rates reached ten percent, a value not seen since the 1980s.

These events have put modern quantitative analysis of the effects of business-cycles on labour markets in the spotlight once again. The first generation of real business cycle (RBC) models that became popular in the 1980s were, for the most part, full-employment models as far as the labour market was concerned: wages clear the market so that the aggregate hours firms want to hire correspond exactly to those that households are willing to offer at the market-clearing wage rate. In such a setting the only two labour market variables the model has implications for are the real wage rate and total hours worked. Not only did these early models found it difficult to match the business-cycle properties of total hors worked, but importantly, they were necessarily silent about such crucial concepts as unemployment and vacancies, as they were not part of the model.

While RBC models, for parsimony's sake, ignored unemployment, early work on labour market models that emphasized search and matching frictions between workers and firms had already begun with the seminal contributions of Mortensen (1970), Diamond (1982), and Pissarides (1985). These early models generated equilibrium unemployment, but they were not business cycle models. It was not until the work of Merz (1995) and Andolfatto (1996) that these two strands of literature were put together and the implications of economy-wide fluctuations for the labour market could be better understood.

I will refer to this framework, where productivity fluctuations are the main driver of the labour market where workers search for jobs and firms search for workers (and are eventually matched, or not) as the Diamond-Mortensen-Pissarides (DMP) model. The goal of this article is to provide a macroeconomic analysis of the Portuguese labour market in the context of other advanced economies in the Organization for Economic Co-operation and Development (OECD) through the lens of this framework.

Starting in the late 1990s the DMP framework has become the workhorse of macro-labour research. And search and matching models of the labour market have been embedded in much more complicated underlying models with numerous frictions and a role for the monetary authority. One issue that has become particularly salient in this line of research is this framework's lack of ability to generate enough volatility in labour market variables. Shimer (2005) documents how the model, when calibrated to deliver the kind of economic cycles one observes in U.S. data generates fluctuations in unemployment and vacancies that are below the ones observed by an order of magnitude. This result became known as the "volatility puzzle" and spurred on a whole literature that attempts to reconcile the model and data. This article shows that this puzzle applies not only to U.S. data but also to OECD, and in particular to Portuguese data also.

There are three dimensions that make the Portuguese labour market stand out from the OECD crowd. The first one concerns worker flows in and out of unemployment. By worker flows I mean the probability that an unemployed worker finds a job in given period of time (the job-finding rate) and the probability that an employed worker loses a job and transits to unemployment (the separation rate). Up until 2011, estimates of these flows obtained through the Instituto Nacional de Estatística's (*INE*) Inquérito ao Emprego (IE) were among the lowest in the OECD. In the sample used here, Portugal has the lowest separation rate and the second lowest job-finding rate, as Table 1 shows (where f represents the job-finding rate and s the separation rate). These rates are estimated by Hobijn and Sahin (2009) and Elsby *et al.* (2011) using data on the number of people unemployed, employed, as well as on unemployment duration for various OECD countries.

In 2011, the *INE* changed the methodology surrounding the survey associated with the IE and these transition rates roughly doubled.[3] Because the sample in Elsby *et al.* (2011) ends in 2007, the flows' estimates found there for Portugal are substantially smaller than the ones found with the new methodology.

Table 1 • Monthly job-finding and separation rates | Percent

	f	s
Australia	20.4	1.7
Austria	15.6	0.8
Canada	23.0	2.4
Czech. Rep	8.1	0.9
Finland	13.4	1.4
Germany	5.8	0.5
Japan	17.2	0.6
Norway	32.0	1.6
Poland	7.2	1.0
Portugal	6.1	0.4
Spain	6.1	1.1
U.K.	13.0	1.0
U.S.	43.2	3.5

Sources: Hobijn and Sahin (2009) and Elsby *et al.* (2012).

The second dimension where Portugal stands out from the OECD norm is in how the economic cycle (measured by output per worker) moves with the labour market. While the statistical correlation between productivity and unemployment (vacancies) is significantly negative (positive) in most countries, it is very close to zero in Portugal (along with a few other countries). It is as if, by virtue of idiosyncratic institutions, the labour market is insulated from the business-cycle at large. What these exact institutions are is a matter for future research to find out, but unions, alternative contractual arrangements, the prevalence of internships, to name just a few, come to mind as possible candidates.

Finally, the third dimension, and one that does not pertain directly to the labour market but exerts on it direct influence, is that the persistency of output per worker is small relatively to other OECD countries. While this finding may seem inconsequential, it does have important implications in the context of the large literature that spawned in response to the volatility puzzle as we will see in section "Using cross-country data to evaluate solutions to the volatility puzzle". Out of all the attempts to reconcile the DMP model with the observed volatility in U.S. labour markets, Hagedorn and Manovskii (2008), henceforth HM, has probably received most of the attention. It proposes a modified version of Shimer (2005) and a different way of calibrating, of disciplining the model using the data, which delivers the kind of volatility we see in U.S data. The present paper shows this strategy does not work when countries (like Portugal) are characterized by productivity processes that have sufficiently low persistency.

In the next sections, I will first review the OECD data in more detail and in particular try to position Portugal's macroeconomic labour market variables in the context of its peer OECD countries. I will then briefly present the class of DMP models I alluded to before in a simple way, and discuss the volatility puzzle. I will finish by showing how the HM solution does not quite work for countries like Portugal that lack sufficiently high persistence in their productivity per worker.

Some cyclical properties of OECD labour markets

The data underlying all results in this article come from unbalanced data panels at a quarterly frequency on vacancies, unemployment, employment, labour force, and real GDP (all in levels) for a set of 16 OECD countries. The proximate sources are the OECD's Economic Outlook Database, the IMF's International Finance Statistics, Ohanian and Raffo (2012), as well as some direct national sources.[4]

While the data collection process for most of these variables is fairly standard across OECD countries, the same cannot be said for the vacancy data. The OECD compiles its vacancy data from a variety of national sources with no harmonized reporting procedures.[5] Nonetheless, to the extent that the majority of data collection differences manifest themselves at low frequencies, the fact that I remove a trend component from the data, keeping only its cyclical part, should help make the vacancy data more comparable across countries.[6]

A number of facts emerge, some new to the literature, some already known, that should provide useful benchmarks for business-cycle models of the labour market. The first finding is that there is substantial variation in the degree of correlation between productivity and unemployment and between productivity and vacancies as shown in Chart 1. Economists use statistical correlations between variables to understand and measure co-movements between them. Correlations can go

from minus one, if the two variables always move in a different direction proportionately, to zero, if the variables are independent from each other, to plus one, if they move in the same direction, proportionately.

Note that most of the correlations are of the expected sign: negative for unemployment, meaning that when the economy is doing well and output per worker is high, unemployment is low and vice-versa; and positive for vacancies, meaning that when output per worker is high vacancies are also peaking, reflecting an increase in the demand for labour on the part of firms. Nonetheless, there are exceptions lying outside the North-West quadrant of the chart. In Spain, for example, it seems like productivity and vacancies do not co-move at all, while productivity and unemployment exhibit a puzzling positive correlation. In countries like Portugal, Norway, Poland, and even Australia, the correlations are very close to zero, suggesting the labour markets there are largely insulated from business-cycle fluctuations. While it is hard to know, at this level of analysis, what is behind this phenomenon, a possibility is that institutions, particular to some countries, may be creating frictions that hinder the transmission mechanism from the business-cycle at-large to the labour market. While studying exactly what those institutions might be is beyond the scope of this piece, one can speculate on possible candidates. The idea here is that anything that impedes the incentives the product market is transmitting from reaching the labour market (and thus slows labour market churning down) may be a candidate. Institutions like unions may trade-off wage growth for employment stability, which would make vacancies and unemployment become less sensitive to changes in the business cycle. Alternative contractual arrangements, like shorter workweeks that allow for total hours to vary while employment is more stable, could produce the same result. Finally, professional internships (state-sponsored or not) can sometimes mask real unemployment (depending on how such internships are registered for unemployment purposes) and could also explain this result.

Table 2 details cross-correlations between variables across time. Taking Portugal and unemployment as an example, the way to read this table is the following: at (x) we read contemporaneous correlations, meaning the correlation between productivity and unemployment period-by-period is -0.082. This is also the value used in Chart 1. But suppose I want to find out the correlation between productivity this quarter and unemployment next quarter, then I read it off of the column labeled x(+1); conversely, if I want to find out the correlation between productivity this quarter and unemployment two quarters ago, I read it off of the column labeled x(-2).

Table 2 • Temporal cross-correlations

	x(-5)	x(-4)	x(-3)	x(-2)	x(-1)	x	x(+1)	x(+2)	x(+3)	x(+4)	x(+5)
Australia											
Unemployment	0.559	0.567	0.536	0.423	0.249	0.056	-0.153	-0.294	-0.378	-0.399	-0.364
Vacancies	-0.536	-0,.55	-0.287	-0.138	0.078	0.230	0.376	0.490	0.505	0.521	0.481
Austria											
Unemployment	0.244	0.228	0.189	-0.019	-0.123	-0.387	-0.480	-0.424	-0.414	-0.322	-0.148
Vacancies	-0.158	-0.079	0.029	0.169	0.333	0.480	0.539	0,538	0.471	0.369	0.248
Canada											
Unemployment	0.456	0.347	0.209	0.041	-0.102	-0.247	-0.358	-0.31	-0.446	-0.431	-0.383
Vacancies	-0.381	-0.267	-0.139	0.014	0.167	0.299	0.394	0.457	0.468	0.455	0.411
Czech. Rep											
Unemployment	0.500	0.402	0.204	0.019	-0.230	-0.435	-0.592	-0.671	-0.675	-0.612	-0.492
Vacancies	-0.381	-0.250	-0.001	0.219	0.457	0.631	0.714	0.717	0.672	0.581	0.449
Finland											
Unemployment	0.440	0.367	0.222	0.062	-0.101	-0.282	-0.435	-0.352	-0.560	-0.584	-0.558
Vacancies	-0.407	-0.299	-0.124	0.042	0.224	0.408	0.496	0.572	0.583	0.601	0.546
Germany											
Unemployment	0.179	0.097	-0.007	-0.125	-0.257	-0.376	-0.439	-0.434	-0.392	-0.297	-0.194
Vacancies	-0.093	0.021	0.147	0.258	0.359	0.445	0.482	0.441	0.342	0.244	0.140
Japan											
Unemployment	0.241	0.166	0.029	-0.106	-0.293	-0.461	-0.571	-0.630	-0.611	-0.508	-0.337
Vacancies	-0.169	-0.061	0.086	0.264	0.457	0.612	0.695	0.681	0.561	0.365	0.139
Norway											
Unemployment	0.433	0.378	0.303	0.237	0.075	-0.038	-0.176	-0.293	-0.342	-0.347	-0.388
Vacancies	-0.514	-0.477	-0.375	-0.205	-0.078	0.056	0.171	0.269	0.327	0.394	0.430
Poland											
Unemployment	0.197	0.279	0.350	0.352	0.312	0.244	0.125	0.049	-0.007	-0.032	-0.007
Vacancies	-0.368	-0.279	-0.141	-0.025	0.113	0.271	0.264	0.214	0.137	0.041	-0.060
Portugal											
Unemployment	0.205	0.170	0.148	0.080	-0.028	-0.082	-0.168	-0.213	-0.260	-0.246	-0.256
Vacancies	-0.135	-0.043	0.042	0.119	0.268	0.262	0.260	0.219	0.161	0.146	0.169
Spain											
Unemployment	0.313	0.399	0.433	0.477	0.477	0.472	0.420	0.369	0.323	0.261	0.203
Vacancies	-0.221	-0.214	-0.154	-0.143	-0.114	-0.076	-0.090	0.001	0.076	0.109	0.146
U.K.											
Unemployment	0.711	0.647	0.487	0.283	0.046	-0.185	-0.392	-0.521	-0.584	-0.567	-0.508
Vacancies	-0.412	-0.270	-0.060	0.181	0.418	0.625	0.741	0.747	0.661	0.551	0.406
U.S.											
Unemployment	0.569	0.550	0.460	0.284	0.041	-0.242	-0.425	-0.536	-0.544	-0.485	-0.393
Vacancies	-0.533	-0.465	-0.329	-0.107	0.157	0.408	0.555	0.608	0.576	0.497	0.404

Sources: OECD and author's calculations.

A quick look at the table shows that for the majority of countries, a picture arises where unemployment reaches its trough roughly three quarters after productivity peaks, which is what happens in Portugal. Vacancies peak roughly two quarters after productivity, unlike what happens in

Portugal where productivity and vacancies peak together. More importantly perhaps, the correlations between labour market variables and productivity remain very low in Portugal throughout the cycle, never exceeding 0.3, while they can reach at least double that value in some other countries like the U.S or the U.K., confirming the view that the business-cycle exerts relatively little influence in the Portuguese labour market.

On the light of this evidence it seems that a model of the labour market that is mainly driven by productivity shocks, where unemployment is low when productivity is high and vacancies move with productivity, may be a bad idea. After all, Chart 1 shows that some countries' labour markets are largely insulated from the business cycle, while others, like Spain, show the opposite behavior, suggesting other mechanisms or institutions may be at work. Yet, this is not the case for all countries. Note also that the close linear relationship between the two sets of correlations in Chart 1 suggests that whatever is driving a wedge between the behavior of productivity and labour market variables affects unemployment and vacancies equally. That is, countries that have a high (absolute) correlation between productivity and unemployment tend to exhibit a high correlation between productivity and vacancies. This suggests that a model where movements in productivity are the underlying force, but there may be some country-specific frictions that vary in strength, may be appropriate.

Charts 2 and 3 further suggest that productivity may be the right driver if one were to build a macroeconomic model of the labour market: there is a fairly strong positive cross-country correlation between the volatility of productivity and that of both unemployment and vacancies. This means that unemployment and vacancies tend to vary more in countries where output per worker varies more. In fact, both unemployment and vacancies are over ten times more volatile than productivity as measured by their standard deviations, and Portugal stands right in the middle of the crowd in this dimension.

Chart 1 • Correlations between productivity and unemployment and between productivity and vacancies

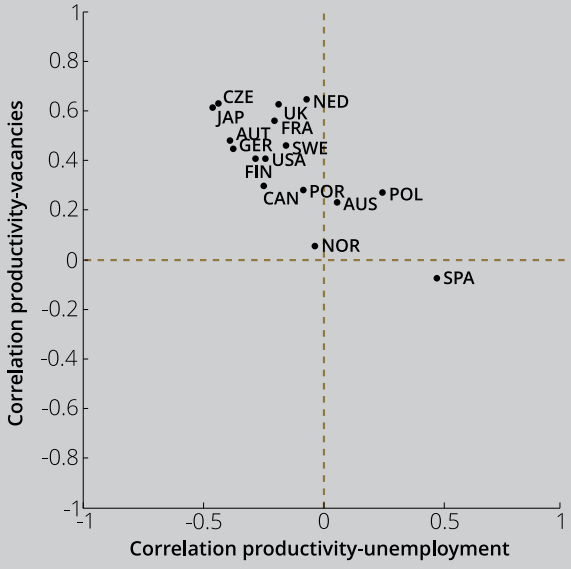

Sources: OECD and author's calculations.

Chart 2 • Volatility of unemployment and productivity | Standard deviation

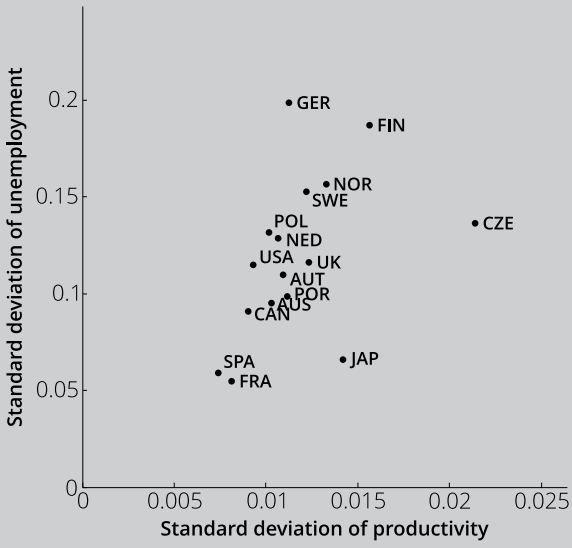

Sources: OECD and author's calculations.

Another important feature of the data is that both vacancies and unemployment are fairly persistent, as given by the correlation between the current quarter and last quarter shown in Chart 4 In this dimension Portugal is very close to the OECD median country, where these labour market variables change substantially but do so slowly, in a persistent way, instead of jumping around. In terms of building a model that features this persistence it is either the case that the productivity that drives the business cycle is very persistent, or the model features some internal mechanism that makes the labour market variables sticky and does not allow them to adjust substantially when productivity changes.

It was with these sorts of relationships in mind that Merz (1995) and Andolfatto (1996) first developed a model connecting the business-cycle with a meaningful model of the labour market featuring equilibrium unemployment and merged the two literatures we referred in the introduction. In the next section I present a simplified version of this work, that I will call the DMP model, focusing more on the labour market than on other features of the economy, and where the main driving engine are productivity shocks, that is based on Shimer (2005).

A productivity-driven model of the labour market

In this model economy there are workers and firms. They may form matches on a one-to-one basis.[7] Workers can thus be employed if they are matched with a firm or unemployed and searching for a job otherwise. Firms can be matched with a worker and producing output or searching for one. Unemployed workers receive a fixed unemployment subsidy and employed workers receive a wage. Firms that are looking for a worker do so by posting a vacancy, that costs them a fixed amount, while firms that are already matched with a worker pocket the profit: the difference between the output they produce and the wage they pay the worker.

Chart 3 • Volatility of unemployment and vacancies | Standard deviation

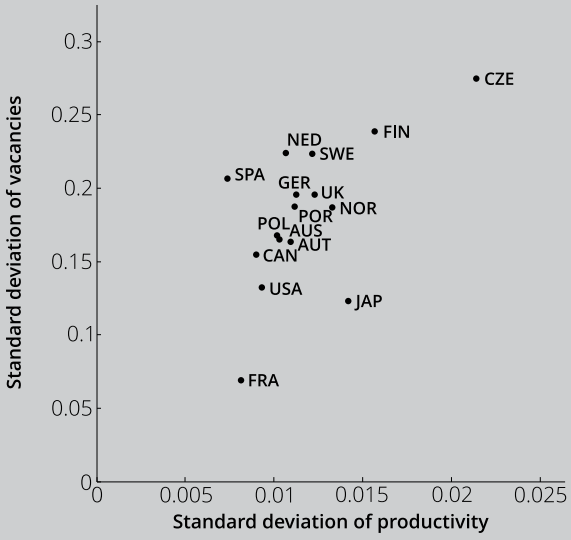

Sources: OECD and author's calculations.

Chart 4 • Auto-correlations for unemployment and vacancies

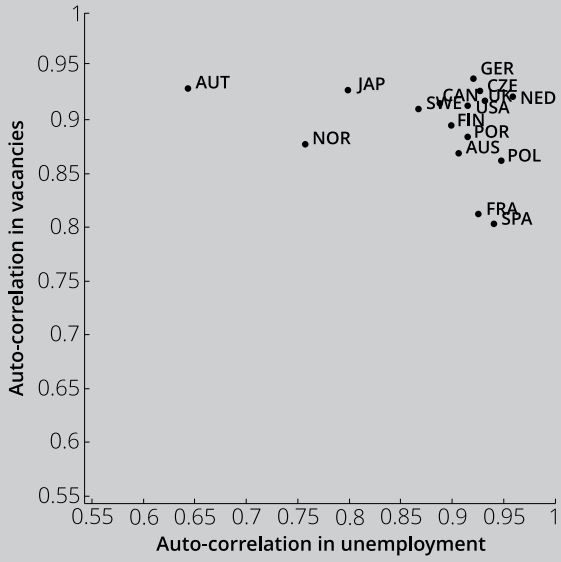

Sources: OECD and author's calculations.

Whether a match is formed or not depends on the aggregate conditions in the labour market. The more unemployed workers there are, the more likely it is that a match is formed; also, the more vacancies firms post, the more likely a match is to be formed. If no vacancies are posted, or there are no unemployed workers, no matches are formed. Once a match is formed it will continue to last until it breaks up which happens according to some fixed probability.

In this economy, the output a matched firm and worker produce is the model analogue of what we called productivity per worker in the last section when talking about the data. Output follows a stochastic process, meaning it is subject to random shocks. In fact, today's output is equal to a fraction of yesterday's output plus a random component that may be positive or negative. This means that output is persistent (because it's partly determined by what past output was), but varies because of the contemporaneous shock.

The wage the worker receives is bargained between the worker and the firm. The exact terms of the bargaining depend on the parties'bargaining power as well as their outside option (zero for the firm and the unemployment subsidy for the worker). Firms take this bargaining into account when trying to decide whether to put up a vacancy or not. In fact, they use all the information available: the probability with which they will find a worker, how much output they are expected to make, the probability that the match breaks up, etc., to compute the present expected gain from posting a vacancy. If the difference between that value and the cost of posting vacancy is positive they decide to post the vacancy.

A crucial variable in this economy is the vacancy-to-unemployment ratio, also known as market tightness. A low ratio means that the market is loose; there are a lot of unemployed workers who find it hard to find jobs, while firms fill their posted vacancies easily without the need to raise wages a lot. A high ratio means that there are a lot of vacancies out there that are not getting filled and unemployment is low. As firms post more vacancies, they need to increase wages to be able to attract workers. This eats into their expected profits, and eventually determines the equilibrium in this model: firms will post vacancies until their expected profit is driven to zero. After that point firms no longer have an incentive to post.

Underlying all this are the productivity shocks that determine how much a match produces. If the shock is good this quarter, firms know that the matches will produce a lot and recall that because productivity is persistent, it is likely that this state of affairs will continue in future quarters. As such, their expected profits will increase, so more vacancies will be posted and more matches will be formed, decreasing unemployment. This is the genesis of the correlations we talked about in the previous sections: as productivity goes up, vacancies go up and unemployment goes down. Productivity and vacancies exhibit a positive correlation, while productivity's correlation with unemployment is negative.

The model was, of course, built to deliver these kinds of relationships, but more than that, we are interested in understanding whether the model can deliver the same magnitudes as seen in the data. To do that properly, and to prevent us from getting whatever result we want, we need to discipline the model. We do that by setting the model's parameters (like the amount of the unemployment subsidy or the probability that a job ends) to match their data counterparts for the countries in our sample. Through this calibration we constrain the simulated data generated by the model to resemble actual data in particular dimensions, while leaving the data dimensions we are interested in, unconstrained.

Charts 5 and 6 show the model's performance in capturing how persistent labour market variables are. In the horizontal axis we have the persistency in the data while the model's persistency is in the vertical axis. The closer the dots are to the 45 degree line (where data and model coincide)

the better the model's performance. The model does a reasonable job capturing the persistence in unemployment, but largely fails in capturing the persistency in vacancies (all observations are below and quite distant from the 45 degree line). While in the data, the correlation between vacancies this quarter and the past quarter in Portugal is roughly 0.9, in the model it is only half that value. This shortcoming is well known in the literature and can be dealt with by introducing mechanisms that slow the adjustment of vacancies down, such as adjustment costs.

Chart 7 shows the model's performance in capturing how volatile labour market variables are relative to the volatility in productivity (which is fixed to be the same in the model and data). So while vacancies in Portugal are roughly 16 times more volatility than productivity, the model suggests they are roughly as volatile. The model's inability to replicate how volatile labour market variables are holds for all countries, that is, the volatility puzzle Shimer (2005) uncovered for the U.S is actually ubiquitous in the OECD.

Using cross-country data to evaluate solutions to the volatility puzzle

There is a voluminous literature dedicated to potential solutions to the volatility puzzle, therefore it is important to be able to evaluate and distinguish between different proposals. The sort of cross-country data this article presents is one possible dimension along which one can scrutinize these alternatives. Here we exemplify this procedure by looking at the solution proposed in HM.

While mostly maintaining the structure of the model, HM proposes that one should target different moments of the data when setting the model's parameters. It notes that the reason standard

Chart 5 • Unemployment auto-correlations: model versus data

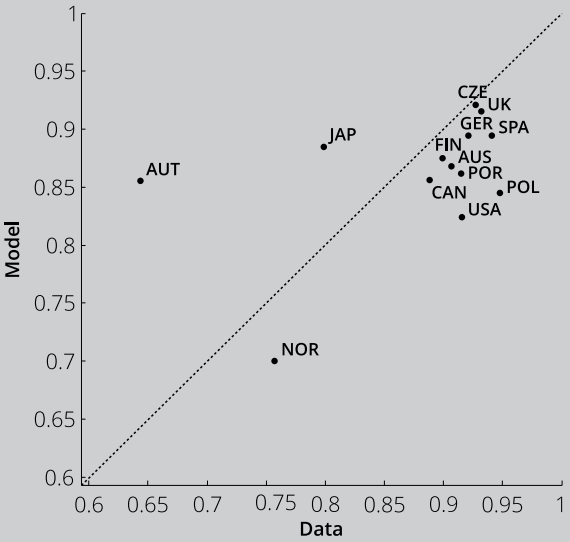

Sources: OECD and author's calculations.

Chart 6 • Vacancies auto-correlations: model versus data

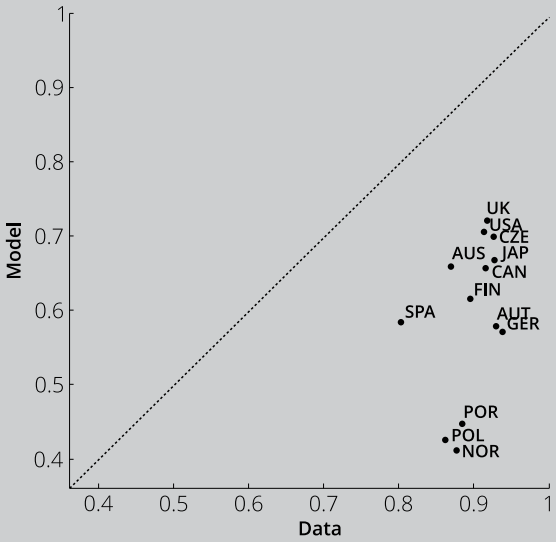

Sources: OECD and author's calculations.

DMP models cannot match the volatility of labour market variables is that, in the model, as the total match surplus varies over the cycle, wages are absorbing much of that variation, while profits vary very little. Recall that it is based on changes in profits that firms set their vacancies, therefore, as profits vary little, so do vacancies, and as a consequence, unemployment.

In the data, by contrast, wages vary a lot less and therefore profits vary more. While part of the literature reacted to this by developing models that emphasized stickier wages, HM took another route and directly targeted the elasticity of wages to productivity. It sets model parameters so that the model replicates the amount of variation in wages one sees in the data. Note that this by itself does not guarantee that the vacancies and unemployment will vary as much as in the data. Even if you set the exact variation in profits that you see in the data, the transmission between profits and labour market variables is still endogenous to the model.

The HM proposal succeeds in bringing the model closer to the data in terms of volatility of labour market variables for most countries, as Chart 8 shows. Yet, for countries like Portugal (and Spain), this attempt is largely unsuccessful and the volatility in the model is still an order of magnitude smaller than in the data. Why is this the case? There are two reasons. The first has to do with how persistent productivity is. In Portugal this persistence is very small, which means that when a good shock occurs firms are less sure that it will last, and therefore they will not raise vacancies by as much as they otherwise would, which results in less volatility in vacancies and unemployment. The second reason has to do with the low job finding rates that characterize the Portuguese (and Spanish) labour markets and it is very straightforward. Conditional on a productivity shock and a given number of posted vacancies, unemployment will decrease by less in an economy where the job-finding rate is smaller, as fewer workers will be able to find jobs.

Chart 7 • Volatility in the labour market: model *versus* data

| Standard deviation of vacancies and unemployment divided by standard deviation of productivity

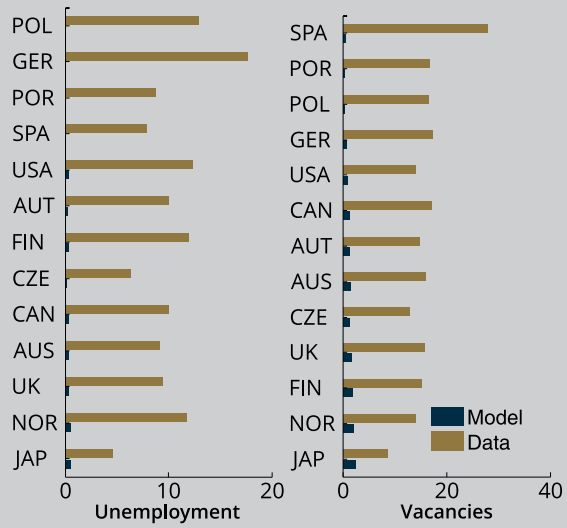

Sources: OECD and author's calculations.

Chart 8 • Volatility in the labour market: model (HM) *versus* data

| Standard deviation of vacancies and unemployment divided by standard deviation of productivity

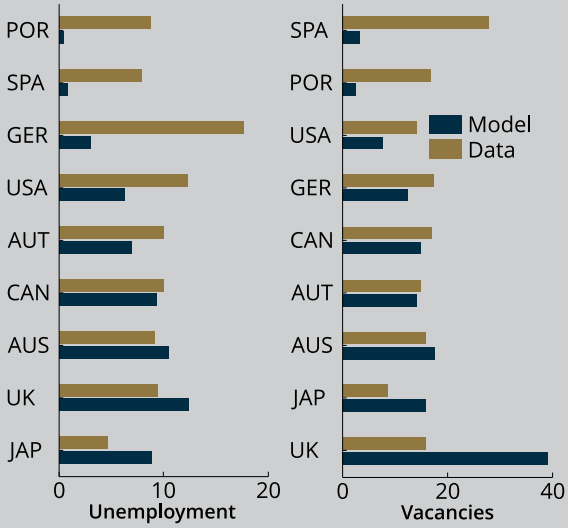

Sources: OECD and author's calculations.

One should note that this last finding in Portugal is subject to revisions in the job-finding rate estimates that resulted from the sampling changes the *INE* introduced in the IE. Nonetheless, Amaral and Tasci (2012) show that as long as the persistency in productivity is small enough, the solution proposed in HM fails to work, even in economies where job-finding rates are relatively high.

Finally, I would like to stress the fact that this type of analysis shows how cross-country statistics can be used to distinguish between proposed solutions to the volatility puzzle.

Conclusions

In the OECD context, the Portuguese labour market shows relatively little connection to the economic cycle, as measured by fluctuations in output per worker. This might have to do with institutions, or other frictions, that blunt the incentives the economy-at-large sends to the labour market.

GDP per worker in Portugal shows relatively little persistency and this explains why a common solution proposed in the literature for the volatility puzzle fails to work for the Portuguese labour market.

References

Amaral, P. S., and M. **Tasci**, (2012), "The cyclical behavior of equilibrium unemployment and vacancies across OECD countries", *Working Paper 1236*, Federal Reserve Bank of Cleveland.

Andolfatto, D. (1996), "Business Cycles and Labor-Market Search", *American Economic Review*, 86(1), 112-32.

Banco de Portugal, (2011), "The Portuguese Economy in 2011", *Annual Report*.

Diamond, P. (1982), "Aggregate Demand Management in Search Equilibrium", *Journal of Political Economy*, 90(5), 881-894.

Elsby, M. W., B. **Hobijn**, and A. **Sahin**, (2011), "Unemployment Dynamics in the OECD", Tinbergen Institute Discussion Papers 11-159/3, Tinbergen Institute.

Hagedorn, M., and Y. **Manovskii**, (2008), "The Cyclical Behavior of Cyclical Unemployment and Vacancies Revisited", *American Economic Review*, 98(4), 1692-1706.

Hobijn, B., and A. **Sahin**, (2009), "Job-finding and Separation Rates in the OECD", *Economics Letters*, 104, 107-11.

Merz, M. (1995), "Search in the labor market and the real business cycle", *Journal of Monetary Economics*, 36(2), 269-300.

Mortensen, D. T. (1970), "Job Search, the Duration of Unemployment, and the Phillips Curve", *American Economic Review*, 60(5), 847-62.

Ohanian, L. E., and A. **Raffo**, (2012), "Aggregate hours worked in OECD countries: New measurement and implications for business cycles", *Journal of Monetary Economics*, 59(1), 40-56.

Pissarides, C. (1985),"Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages", *American Economic Review*, 75(4), 676-90.

Shimer, R. (2005), "The Cyclical Behavior of Equilibrium Unemployment and Vacancies", *American Economic Review*, 95(1), 25-49.

Notes

1. The opinions expressed in the article are those of the author and do not necessarily coincide with those of Banco de Portugal or the Eurosystem. Any errors and omissions are the sole responsibility of the author.

2. Banco de Portugal, Economics and Research Department.

3. For more details, please see Box 4.1 of the *Annual Report – The Portuguese Economy* in 2011 (2011) pp.135-138.

4. Please see Amaral and Tasci (2012) for a more technical report on this research and for a detailed description of all the sources.

5. In particular, for Portugal, the vacancies data are collected from the *Instituto do Emprego e da Formação Profissional* that in turn collects it, at a monthly frequency, from the jobs posted by firms at the various Employment Centers across the country.

6. All variables are in logs and are detrended using the Hodrick-Prescott filter with a smoothing parameter of 1600.

7. Please see Amaral and Tasci (2012) for a more detailed description of the model.