

Natural Unemployment and Activity Rates in Italy: Flow-based Determinants and Implications for Price Dynamics*

Francesco D'Amuri[†] Marta De Philippis[†] Elisa Guglielminetti[†]
Salvatore Lo Bello[†]

October 2022

Abstract

Despite the magnitude and cyclicity of transitions into and out of the labor force, the literature generally considers unemployment as a sufficient statistic of labor market slack. We question this view by jointly estimating natural unemployment and participation rates through a Phillips curve informed by structural labor market flows. Focusing on Italy, a country where flows into and out of the labor force are particularly large, we find that the participation margin accounts for a significant share of total slack and explains one third of the missing inflation that followed the 2011 Sovereign Debt Crisis. Exploiting a reform that sharply and unexpectedly increased the statutory retirement age and expanded labor supply without directly affecting unemployment, we confirm that neglecting the participation gap biases inflation forecasts.

Keywords: Labor market flows, Labor supply, Demographic trends, Phillips curve, Inflation.

JEL Classification: J11, J21, J64, E32

*Corresponding author: francesco.damuri@bancaditalia.it. We are grateful to Gianni Amisano, Regis Barnichon, Stefano Eusepi, Claudio Michelacci, Marianna Kudlyak, Stefano Neri and Aysegul Sahin and to seminar participants at the Bank of Italy, the Einaudi Institute for Economics and Finance, the European Central Bank, the European Commission (Output Gap Working Group) and the 4th Annual European System of Central Banks Research Cluster 2 workshop for their comments. We also thank Stefano Eusepi, Petteri Juvonen and Meri Obstbaum for kindly sharing their Matlab codes. Usual disclaimers apply. All errors are ours.

[†]Bank of Italy

1 Introduction

The assessment of labor market slack is crucial for understanding an economy’s cyclical position and its price dynamics. There is a long tradition in the macro literature that considers the unemployment rate as the statistic conveying all relevant information on slack. At the same time, it is well-known that flows into and out of the labor force are way larger than those between employment and unemployment [Blanchard and Diamond, 1990]. For instance, in the US, over the 1990-2020 period flows into and out of activity were 2.8 times larger than flows between employment and unemployment.¹ The same is true in the main European countries.² Moreover, it has recently been documented that, besides being very large, these flows account for a relevant share of the cyclical variation in the unemployment and participation rates in the US [Elsby et al., 2015, Hobijn and Şahin, 2021] and other countries. Hence, several recent contributions have analyzed the potential role of the participation margin for price dynamics by constructing new measures of search intensity that consider the whole non-employed population and go beyond the official unemployment rate [Abraham et al., 2020]; this issue also attracted the attention of top Federal Reserve officials: in a 2021 speech, L. Brainard underlined that “changes in labor force participation contain important information about the strength of the labor market that is not captured in the unemployment rate”.

However, to the best of our knowledge, none of the existing papers takes into account that movements into and out of the labor force can affect price dynamics without necessarily being related to the number of job seekers or to search intensity. For example, consider a reduction of the transitions from employment to inactivity due to an increase in statutory retirement age. This rise in labor supply, not directly related to unemployment, could affect wages either negatively, if workers of different ages are perfect substitutes and firms move down on their labor demand curve,³ or positively, for example if imperfect substitution prevails and a reduction in the quit rate of experienced workers results in an increase in demand for younger workers as well.⁴ This implies that, everything else equal, such rise in the participation rate would affect price dynamics and that the direction of the effect is *a priori* ambiguous. Hence, not only studying the evolution of wages and prices requires to take into account the participation margin over and beyond its obvious nexus with the unemployment rate, but it is also necessary to estimate in which direction its movements affect the overall level of labor market slack.

In this paper we tackle this issue by developing a unified framework for the estimation of the natural unemployment and participation rates, which together define the

¹An average of 11.4 million individuals moved each month either into or out of activity, as opposed to 4 million between employment and unemployment

²According to harmonized data recently made available [Eurostat, 2020], the ratio between the sum of quarterly gross flows into and out of activity over the flows between employment and unemployment during the 2010-2020 period was equal to 1.5 in Spain, 2 in France, 2.5 in the United Kingdom and 4.6 in Italy.

³Such conclusion would be put into question if older workers enjoy a stronger bargaining position due to more stringent employment protection or the presence of anti-discrimination clauses.

⁴Exploiting the different exposure of Italian firms to the unexpected reform also analyzed in this paper, two recent papers [Bianchi et al., 2022, Carta et al., 2021] do not find sizeable effects on wages. Nevertheless, as they identify the impact looking at the exogenous variation across firms, these papers are not able to quantify the market level effects of the reform.

labor market condition implying a neutral impact on inflation. To do so, we propose an augmented version of the Phillips curve relationship, which explicitly accounts for the participation margin and is informed by labor market flows. The model allows us to estimate an enriched measure of economic slack that, by taking into account participation gaps, contributes to explaining the low inflation regime of the years following the Great Recession and the Sovereign Debt Crisis.

We apply our framework to Italian data, a choice motivated by two main reasons. First, Italy is one of the countries where fluctuations into and out of the labor force are largest compared to those between employment and unemployment [Eurostat, 2020]. Second, we leverage a large and unexpected increase in full retirement age that took place in 2012 to study how an exogenous shift in participation – unrelated to unemployment or job search intensity – is captured by our Phillips curve framework and helps interpreting price dynamics in recent years.

Our estimation framework builds on Crump et al. [2019] and consists of two main steps. First, we estimate the structural participation and unemployment rates, which are determined by purely structural factors such as changes in the demographic composition and in the labor market environment (preferences, institutions, matching technology). To this aim, we follow Crump et al. [2019] in using a flow-based model of unemployment (and participation) dynamics [Shimer, 2012], but we consider as labor market states not only employment and unemployment but also inactivity. We extract the trend components of labor market flows for six demographic cells defined by gender and age and use them to compute the steady state unemployment and participation rates. We then aggregate the cell-specific rates to obtain the overall structural unemployment and participation rates. In our second step, we estimate the natural unemployment and participation rates, i.e. those consistent with constant inflation, through a forward-looking Phillips Curve that uses as anchors the previously obtained structural rates.

Crump et al. [2019]’s approach is particularly well-suited for the scope of our analysis: by studying labor market flows over time, it allows us to identify and better interpret the determinants of changes in the structural rates. In this respect, we innovate by adding inactivity as a third labor market state, thus obtaining an estimate of the structural participation rate. It is important to notice that the same underlying structural flows jointly determine the unemployment and the participation rates, giving discipline to the exercise. Moreover, we allow for both an unemployment and a participation gap in the estimation of the Phillips curve, retrieving their separate effects on price dynamics. This is a distinctive feature of our work, and constitutes the main innovation of our framework.

We find that the structural unemployment rate in Italy exhibited little fluctuations throughout the period 1984–2018,⁵ whereas the structural participation rate steeply increased by about 6 percentage points in the same period. This suggests that, even though they are determined by the same underlying flows, these rates feature a very different response to cyclical and structural shocks. Our framework allows us to dig deeper into the

⁵We start our analysis in 1984 due to data availability. The last year included in the analysis is 2018 since starting with 2019 a temporary pension reform (so called Quota 100), reduced pension eligibility requirements and might interfere with our estimates of the structural and natural rates. In a complementary paper we study the effect of Covid-19 on the unemployment and participation gaps [D’Amuri et al., 2022].

factors shaping these dynamics. We show that most of the rise in the structural activity rate was due to the increasing participation of older workers (55-64); in turn, this can be traced back to a marked decline in the flows from employment to inactivity, very likely linked to a number of pension reforms that took place in Italy during the time period of our analysis.

We then turn to the Phillips curve estimation to assess the role of the participation gap in determining labor market slack and in shaping price dynamics. We find that the participation gap provides autonomous and relevant information on price developments, accounting for about 70% of the contribution provided by labor market slack. Its role becomes especially relevant over the years following the double-dip recession, when the participation margin alone explains one third of subdued inflation compared to its long-run trend, thus helping in explaining the missing inflation puzzle. Given the peculiarity of the Covid-19 crisis, in this paper we do not discuss the role of the participation margin for price dynamics during the health emergency, which we analyze in a separate contribution [D'Amuri et al., 2022].

To provide further evidence on the effects of participation on price dynamics, we exploit our rich framework to evaluate the impact of a far reaching pension reform that took place in Italy in 2012 (the Fornero reform). The reform, which significantly increased the statutory retirement age, was unexpected and swiftly implemented, leaving no room for firms and workers to adapt turnover and labor supply decisions in advance. We first evaluate the impact of the reform by deriving counterfactual structural participation and unemployment rates and comparing our best forecast for the 2011-2015 period based on the pre-reform trends with new estimates using all the post reform data up to 2015. We find that the reform increased structural participation of older individuals (55-64) – through a sharp contraction of the employment to inactivity transitions – and had negligible effects on the participation of the other age classes, implying a 0.7 p.p. increase in the aggregate structural activity rate. Importantly, the reform had virtually no impact on the structural unemployment rate. Finally, we use our Phillips curve model to derive the evolution of the natural unemployment and participation rates by conditioning on the estimated counterfactual paths of the structural rates only. We find a large effect of the pension reform on the natural participation rate (+0.7 p.p.) and no noticeable effect on the natural unemployment rate, showing that there is no clear trade-off between increasing participation of the elderly and unemployment. An expansion of the natural participation rate also implies that, following the reform, a given observed participation rate is associated with lower inflation pressures. Failing to account for changes in natural participation would have thus resulted in a positive bias of inflation forecasts after the reform.

Our work relates to the large literature on the flow-based analysis of labor market dynamics (Choi et al. [2015], Crump et al. [2019], Elsby et al. [2019], Fujita and Ramey [2009], Gomes [2012], Petrongolo and Pissarides [2008], Shimer [2012], among others).⁶ In particular, we extend the framework of Crump et al. [2019] by explicitly taking into account the participation margin. We also contribute to the literature that studies the

⁶For an application of the Shimer [2012] approach to Italy see Sestito [1988] and Rosolia [2014].

role of labor supply fluctuations on cyclical dynamics (Elsby et al. [2015], Garibaldi and Wasmer [2005], Krusell et al. [2017, 2020], King [2011], Kudlyak and Schwartzman [2012], Lalé [2013], Pries and Rogerson [2009], Strand and Dernburg [1964]) and to the literature proposing alternative measures of labor market slack (Aaronson et al. [2014], Abraham et al. [2020], Bell and Blanchflower [2013], Faberman et al. [2020], Gordon [2013], Hornstein et al. [2020], among others). Relative to these papers, we are the first to gauge the relative importance of the unemployment and the participation gaps in the context of the Phillips curve estimation, exploiting a consistent and comprehensive framework.

We also contribute to the extensive literature on Phillips curve estimation (see Ball and Mazumder [2019], Coibion and Gorodnichenko [2015], Del Negro et al. [2020], Stock and Watson [2020] for recent discussions) by showing that the inclusion of the participation gap contributes to explaining price dynamics, overcoming some failures of the standard Phillips curve after the double-dip recession (Ball and Mazumder [2011], Bobeica and Jarociński [2019]). Our approach lies in the tradition of estimating reduced-form Phillips curve models, without attempting to solve structural identification issues such as those addressed by Barnichon and Mesters [2020], Mavroeidis et al. [2014], and McLeay and Tenreyro [2019].

On the theoretical side, models considering the participation margin within a traditional New Keynesian framework include Campolmi and Gnocchi [2016], Erceg and Levin [2014] and Galí et al. [2011]. We view our work as complementary to these papers, as none of them jointly estimates the natural unemployment and participation rates embedded in the model-based Phillips curve. Moreover, our work provides important insights also for search and matching models, which usually do not model the participation margin (few exceptions are Cairó et al. [2022] and Krusell et al. [2017]). We document that assuming fixed labor supply biases the estimation of the matching function and of matching efficiency (as also pointed out by Barnichon and Figura [2015] and Hall and Schulhofer-Wohl [2018]) and that this bias varies along the cycle, therefore changing the interpretation of movements in the Beveridge curve, for instance. Finally, our estimates of the effects of a pension reform that increased retirement age are informative for the debate about the macroeconomic effects of ageing (see for instance Acemoglu and Restrepo [2017], Barnichon and Mesters [2018], Engbom [2019], Feyrer [2007]).

The rest of the paper is organized as follows. Section 2 defines the main quantities at study; Section 3 presents the methodology to obtain structural activity and unemployment and outlines the results. Section 4 sets up a Phillips curve model to estimate natural participation and unemployment rates and to understand their role for price dynamics; using such framework, Section 5 analyzes the impact of the 2012 pension reform on structural and natural rates. Finally, Section 6 concludes.

2 Definitions

We distinguish two concepts that have often been used interchangeably in the previous literature: i) the *structural* (or trend) unemployment and participation rates, derived by extracting the trend components from labor market flows and evaluating them at

the steady state; these are determined by purely structural factors like changes in the demographic composition and in the labor market environment (preferences, institutions, matching technology); and ii) the *natural* unemployment and participation rates, that are the rates consistent with constant inflation, estimated within the context of a Phillips curve framework that uses price and wage dynamics to infer the degree of slack in the economy.⁷

As in Crump et al. [2019], we adopt the following decomposition of the unemployment rate:

$$u_t = \bar{u}_t + \underbrace{(u_t - u_t^*)}_{x_t^u} + \underbrace{(u_t^* - \bar{u}_t)}_{z_t^u}, \quad (1)$$

where u_t is the actual unemployment rate; \bar{u}_t is the structural (or trend) unemployment rate, u_t^* is the natural unemployment rate and $x_t^u = (u_t - u_t^*)$ is the unemployment gap. $z_t^u = (u_t^* - \bar{u}_t)$ is the gap between the natural and the structural unemployment rates. While the structural unemployment rate tracks the evolution of unemployment due to structural forces, the natural unemployment rate connects the real and the nominal side of the economy. We use \bar{u} to discipline u^* , as we assume that u^* converges to \bar{u} in the long-run; in the short-run, monetary policy shocks or temporary changes in price or wage setting may affect u^* without having an impact on \bar{u} . For example, the introduction of an ex-post wage indexation mechanism in an economy with accelerating inflation would temporarily drive u^* up but would have no effect on structural unemployment.

We introduce a similar decomposition also for the participation rate:

$$p_t = \bar{p}_t + \underbrace{(p_t - p_t^*)}_{x_t^p} + \underbrace{(p_t^* - \bar{p}_t)}_{z_t^p}, \quad (2)$$

where p_t is the actual participation rate; \bar{p}_t is the structural (or trend) participation rate; p_t^* is the natural participation rate, consistent with constant inflation. $x_t^p = (p_t - p_t^*)$ is the participation gap; for a given unemployment gap, this variable conveys information on the additional degree of slack in the economy. $z_t^p = (p_t^* - \bar{p}_t)$ is the gap between the natural and the structural participation rates. Again, to discipline the estimation of p_t^* , we assume that p_t^* will converge to \bar{p}_t in the long-run (i.e. that z_t^p will converge to zero). Temporary shocks can move p^* without affecting \bar{p} , hence determining a positive or negative gap z_t^p . For instance, temporary tax cuts meant to favor women re-employment after maternity leave would raise p^* and trickle down to structural labor force participation only insofar they trigger a permanent rise in labor force participation.

3 Structural unemployment and activity rates

In this section we describe the methodology used to estimate the structural unemployment and activity rates, which we then use as anchors for the estimation of the natural rates described in Section 4.

⁷See also Crump et al. [2020] for a discussion about various concepts of unemployment rate benchmarks.

3.1 Estimation

The estimation procedure involves four steps (for a full description with all the technical details see Section A of the Appendix). First, we use the Italian Labor Force Survey micro data to estimate labor market flows between the three labor market states (employment E , unemployment U and inactivity N) over the 1984-2018 period for six demographic groups defined by three age classes (15–34, 35–54 and 55–64) and gender.⁸ Following the existing literature (for instance Barnichon and Mesters [2018], Elsby et al. [2015] and Shimer [2012]), we perform two important adjustments to the measured flows: i) we make them consistent with the dynamics of the stocks (*margin error correction* - MEC) and ii) we derive continuous-time flow rates from discrete time transition probabilities, in order to account for the possibility of multiple transitions taking place within the observation window (*temporal aggregation correction* - TAC). In this way, we obtain six hazard rates for each group g :

$$\{f_{g,t}^{NU}, f_{g,t}^{NE}, f_{g,t}^{EU}, f_{g,t}^{EN}, f_{g,t}^{UE}, f_{g,t}^{UN}\}_{t=1984q1}^{2018q4},$$

where $f_{g,t}^{XY}$ is the transition rate between labor market state X and Y for demographic group g at time t .

Second, following Tasci [2012], we decompose each labor market flow (in each demographic cell) into a stochastic trend and a stationary cyclical component, using an unobserved component model which takes into account their joint dynamics with real log GDP (see Section A of the Appendix for details).⁹ The outcome of this second step are the trend components of the flow rates ($\bar{f}_{g,t}^{XY}$), which represent the building blocks of the structural unemployment and activity rates computed in the next step.

Third, to compute the structural rates of unemployment and participation, we rely on the notion of flow-consistent rates (as in Shimer, 2012). Let $U_{g,t}$, $E_{g,t}$ and $N_{g,t}$ be the relevant stocks of unemployment, employment and inactive population for demographic group g at time t . The evolution of the stocks over time depends on the hazard rates through the following differential equations:

$$\dot{U}_{g,t} = f_{g,t}^{EU} E_{g,t} + f_{g,t}^{NU} N_{g,t} - (f_{g,t}^{UE} + f_{g,t}^{UN}) U_{g,t}, \quad (3)$$

$$\dot{E}_{g,t} = f_{g,t}^{UE} U_{g,t} + f_{g,t}^{NE} N_{g,t} - (f_{g,t}^{EU} + f_{g,t}^{EN}) E_{g,t}, \quad (4)$$

$$\dot{N}_{g,t} = f_{g,t}^{UN} U_{g,t} + f_{g,t}^{EN} E_{g,t} - (f_{g,t}^{NU} + f_{g,t}^{NE}) N_{g,t}. \quad (5)$$

⁸The choice of the starting and ending point is dictated by data availability. Instead, the choice of the age groups is dictated by the observed patterns of participation, which are increasing over the age 15–34, substantially flat between 35 and 54, and progressively declining in the region 55–64 (Figures C.1 and C.2). Throughout the paper, we will therefore estimate aggregate rates for the population 15–64, as the overwhelming majority of the changes in the participation patterns occur before age 64 (see again Figures C.1 and C.2).

⁹Differently from Crump et al. [2019], we employ data on GDP in order to estimate the structural components of labor market flows mainly because we are constrained by the length of our time series. In particular, our time series goes from 1984 to 2018 and encompasses only five recessionary episodes in Italy, some of them very short-lived; Crump et al. (2019) rely instead on a longer time series (from 1960 to 2018) that encompasses eight US recessions according to the NBER definition. Moreover, letting the cyclical component of flows respond to cyclical fluctuations is consistent with standard search and matching models [Mortensen and Pissarides, 1994], in which labor market flows change over the business cycle.

Under the assumption of constant transition rates, we can use (3), (4) and (5) to solve for the steady-state levels of U_g^* , E_g^* and N_g^* , by setting $\dot{U}_{g,t} = \dot{E}_{g,t} = \dot{N}_{g,t} = 0$. In order to obtain the trend unemployment and participation rates, we evaluate equations (3)-(5) in steady state plugging in the estimated trend components of the flows, $\bar{f}_{g,t}^{XY}$ (as in Tasci [2012] and Crump et al. [2019]). Let us then formally define the structural unemployment and participation rates as:

$$\bar{u}_{g,t} = \frac{\bar{U}_{g,t}^*}{\bar{U}_{g,t}^* + \bar{E}_{g,t}^*}, \quad \bar{p}_{g,t} = \frac{\bar{U}_{g,t}^* + \bar{E}_{g,t}^*}{\bar{U}_{g,t}^* + \bar{E}_{g,t}^* + \bar{N}_{g,t}^*}.$$

Plugging in the equilibrium values, and using the fact that total population is normalized to 1, we can solve for the structural unemployment rate $\bar{u}_{g,t}$ and participation rate $\bar{p}_{g,t}$ of each group g as a function of the structural rates $\bar{f}_{g,t}^{XY}$.¹⁰

Forth, we aggregate structural unemployment and activity rates using a weighted average of the group-specific ones. In particular, for the aggregation of the group-specific participation rates, the weight of group g at time t is represented by its share in the total population at a given point in time. Let us denote this population weight as $\omega_{g,t}^p$, such that $\sum_g \omega_{g,t}^p = 1 \forall t$. Hence, the aggregate structural participation rate is computed using $\omega_{g,t}^p$ as weights:

$$\bar{p}_t = \sum_g \omega_{g,t}^p \bar{p}_{g,t}. \quad (6)$$

The aggregation of group-specific unemployment rates is slightly more involved. In this case the weights are a combination of the group-specific weights in the total population (the $\omega_{g,t}^p$ defined above) and their incidence in the active population. Therefore, the aggregate structural unemployment is calculated as follows:

$$\bar{u}_t = \sum_g \omega_{g,t}^p \underbrace{\frac{\bar{p}_{g,t}}{\bar{p}_t}}_{\bar{\omega}_{g,t}^u} \bar{u}_{g,t}, \quad (7)$$

where $\bar{\omega}_{g,t}^u$ can be interpreted as the structural weight in labor force of group g at time t . It represents the share of structural active population (that is, the active population identified by the structural rates) accounted for by the specific group g . For instance, for a given share of the elderly in the population ($\omega_{g,t}^p$), their weight on the structural unemployment rate will typically be smaller because their structural participation rate is lower than the average one.

¹⁰Note that in this setting we are assuming that the population of each group is constant over time. This is a standard assumption in the literature (which, for the papers that do not explicitly consider inactivity, translates into assuming that the active population is constant over time). Barnichon and Mesters [2018] discuss the possible bias generated by a changing within group population in the US. They show that it is negligible, since flows tend to be much larger than population growth in the US. We find a similar result for Italy: for instance, on average, for all demographic groups, flows from inactivity to employment and vice versa contribute about 10 times more than changes in population to variations in the stock of employed individuals. Indeed, flows are more than 10 times larger on average than the deviation between the growth rates of the population of employed individuals and that of the overall population.

3.2 Results

Structural unemployment rate

We plot the estimated series of the aggregate structural unemployment rate in Figure 1. Notice that the structural unemployment rate was essentially flat between 1984 and 1995, to then assume a slow but continuous downward trend until 2008, when it started to rise again until 2015. Since then, the structural unemployment rate appears stable. Overall, its time series smooths out the large oscillations of the actual unemployment rate over the business cycle. In 2018q4, the aggregate structural unemployment rate is estimated to be at 9.2%. The negative gap between the aggregate unemployment rate and the actual one since the sovereign debt crisis is the result of the joint contribution of all demographic groups (see Figure C.3), meaning that for all groups the actual unemployment rate lies above its structural level.

In order to understand what generates the dynamics between 1984 and 2018, Figure 2 shows the evolution of the two components of the aggregate structural unemployment rate: i) the weights of each demographic group, and ii) the group-specific structural unemployment rates.¹¹ The downward trend started in 1995 was brought about by the fast change in the structure of active population, with the 15–34 age group losing share in favour of the group 35–54, characterized by a substantially lower level of structural unemployment. From 2000 onward, the negative trend was further boosted by the increase in the share of the age group 55–64, also characterized by low trend unemployment. At the same time, the rapid increase in the structural unemployment rate of the youngest groups represented a counteracting force, resulting in a slight increase of the aggregate rate after 2008. From 2015, due to the flattening of the structural unemployment of the youngest groups, the aggregate rate appears stable. Figure C.4 provides a decomposition of the dynamics of the aggregate rate into two components: changes in weights (demographics and structural participation dynamics) and changes in within-group rates. The main takeaway is that population shares provided downward pressure on the structural unemployment rate throughout the period, whereas within-group components pushed the aggregate rate up.

Structural participation rate

As for the structural participation rate, Figure 3 reveals that it constantly increased since 1995, reaching the level of 66.3%, in 2018q4. Overall, it followed closely the evolution of the observed participation rate; however, the structural component kept increasing despite the flattening of the actual rate between 2000 and 2010. Indeed, our filtering technique ascribes the decrease in the actual participation rate of prime-age men and of the youth during that period mainly to cyclical conditions, while identifying the increase in the participation rate of the elderly as structural (see Figure C.8).

¹¹Figure C.5 in the Appendix displays these group-specific structural rates together with the actual unemployment rate in each group.

Figure 4 analyzes the determinants of its evolution over time. It reveals that the increase was initially driven by the rising weight on the population of prime age individuals, characterized by higher activity rates; from the early 2000s it was instead mostly driven by the strong growth in the structural activity rate of the eldest groups, together with their increasing weight in the population.¹² Figure C.7 provides a decomposition of the dynamics of the aggregate rate into two components: changes in weights (demographics) and changes in within-group rates. We find that population dynamics brought about positive pressure on the structural participation rate, but the main determinants of the aggregate trend were within-group dynamics.

To dig into the causes of these strong trends in the group-specific structural activity rates, we decompose them into components due to the underlying flows. We divide the six flows into three groups (exit: EN, UN; entry: NE, NU; churn: EU, UE)¹³, and let only one of them vary over time, fixing the others at their sample mean, as in Shimer [2012]. To fix ideas, when we let vary only the exit margin (EN and UN), we construct a counterfactual series of trend activity, i.e. the one generated by movements in the exit margin only. We find that different forces have driven the trends for the different demographic groups (Figure 5). For the youth, the observed reduction in trend activity was primarily due to a slower entry. For prime age individuals, both the entry and the exit margins played a relevant role: while for men they almost exactly offset each other throughout the period, for women they both contributed to the increasing participation, especially the entry margin. Finally, the exit margin clearly drove almost all the increase in the activity rate of the elderly, with an additional push coming from the entry margin. Focusing on the exit margin, we distinguish the contribution of the EN and the UN flow (Figure 6), concluding that the increase in participation of the elderly was entirely accounted for by the reduction in the EN flow.

Overall, the aggregate dynamics in the structural activity rate reflected primarily the unprecedented increase in the structural activity of the elderly, which – on the basis of the results of our decomposition – can be traced back to the strong decline in their EN flows. Part of the change in the behavior of these groups of workers was arguably due to a number of pension reforms, starting in the 2000's.¹⁴ In Section 5, we focus our attention on the Fornero reform, that took place in 2012, and study its effect on both the structural and the natural rates.

4 Natural unemployment and activity rates

We now turn to the estimation of natural unemployment and activity rates. Our estimation framework leverages an augmented version of the standard Phillips curve that exploits

¹²Again, Figure C.8 in the Appendix displays the group-specific structural activity rates together with the actual rate in each group.

¹³We follow the same groups definition of Elsby et al. [2019].

¹⁴For instance the Amato reform (1992) increased the age requirement in order to obtain the old age pension; the Dini reform (1995) gradually shifted the pension system from a defined benefit to a notional defined contribution system; the Moroni and Prodi reforms (2004 and 2007 respectively) also modified the age requirements needed to claim the full pension benefits.

the information contained in the structural rates estimated in the previous Section.

4.1 The augmented Phillips Curve

The Phillips curve is one of the building blocks of New Keynesian models, which provide a microfoundation to what originally represented just an empirical regularity. In particular, it posits that price and/or wage dynamics depend on three main factors: i) inflation inertia and/or expectations, ii) demand factors, iii) supply factors. The New Keynesian tradition has also emphasized the importance of taking into account inflation expectations since agents are forward-looking.

Our approach will follow that of a standard Phillips curve framework. However, in the standard approach demand factors (i.e. economic slack) are generally captured by the unemployment gap, which is the difference between the observed (u_t) and the natural (u_t^*) unemployment rates,¹⁵ and the participation margin is usually neglected. Instead, we take into consideration also the role of changes in participation, which are likely to provide additional information on labor market slack. Indeed, for a given unemployment gap, a participation rate that is below its natural level could be an indication that the economy is running below potential, as some inactive workers could switch to activity.

While some theoretical papers consider the participation margin within a traditional New Keynesian model (Campolmi and Gnocchi [2016], Erceg and Levin [2014] and Galí et al. [2011]), they do not estimate the model-based Phillips curve. Here we move one step further by proposing an augmented Phillips curve model where price and wage inflation are also related to the participation gap, defined as the difference between the observed participation rate and the unobserved level p_t^* , consistent with constant price inflation.

Differently from Erceg and Levin [2014], who use institutional projections of the unemployment and participation rates as proxies for u^* and p^* , we aim at estimating the unobserved natural unemployment and participation rates. Hence, we extend the approach of Crump et al. [2019], using a rich state-space model that allows us to jointly estimate the Phillips curve, the unemployment gap and the participation gap by making standard assumptions on the data generating process. The estimation also exploits the information from labor market flows using the structural unemployment and participation rates previously estimated (\bar{u}_t, \bar{p}_t) as an anchor for u^* and p^* . Furthermore, our setup builds on the insights of a recent literature highlighting the importance of inflation expectations to explain price dynamics (Ball and Mazumder [2019], Coibion and Gorodnichenko [2015]) and incorporates them to pin down the inflation trend.¹⁶

Our approach lies in the tradition of estimating reduced-form Phillips curve models, without attempting to solve structural identification issues such as those addressed by

¹⁵The natural unemployment rate is defined as the level of unemployment for which price inflation remains stable in absence of supply shocks.

¹⁶Although it is known that the textbook Phillips Curve model failed to account for the missing disinflation during the Great Financial Crisis and for the missing inflation over the ensuing recovery (Ball and Mazumder [2011], Bobeica and Jarociński [2019]), recent contributions show that some refinements of this standard tool considerably improve its performance also in the last years. For instance, Coibion and Gorodnichenko [2015] and Ball and Mazumder [2019] argue that the puzzling inflation dynamics can be explained within the context of the Phillips curve framework when inflation expectations are properly taken into account.

Barnichon and Mesters [2020], Mavroeidis et al. [2014] and McLeay and Tenreyro [2019]. However, by including observable measures of inflation expectations and using external information to anchor the estimates of the natural rates such concerns should be mitigated.

4.1.1 Model specification

We estimate a generalized version of the Phillips Curve which nests many specifications commonly used in the literature and can be microfounded by a New Keynesian model like the one presented in Galì [2011]. We sketch here only the main building blocks of the model, leaving to the Appendix (Section B) the description of the full specification.

Consistent with empirical regularities and the theoretical insights of New Keynesian models, we expect inflation to depend negatively on the unemployment gap, which captures negative demand factors. Recall that we denote the unemployment gap at time t by $x_t^u = u_t - u_t^*$. We augment the standard specification by including an additional indicator of demand pressures, that is the participation gap $x_t^p = p_t - p_t^*$, where p^* is the natural participation rate.¹⁷

Like Crump et al. [2019], we model supply-factors through an AR(1) process $\Delta(a_t)$, that should absorb the variation in prices and wages due, for instance, to productivity growth and cost-push shocks.

Following the New Keynesian tradition, we allow current inflation to depend on future inflation expectations. By assuming that agents form their expectations rationally, this implies that current inflation depends not only on the current state of the economy (as captured by the current unemployment and participation gaps), but also on its expected future realizations. Therefore the model features the present discounted value of expected unemployment and participation gaps. The quarterly discount rate β is fixed at the standard value of 0.99, which implies an annual interest rate of about 4%.

Furthermore, the model admits a time-varying inflation trend (π_t^*) which we jointly estimate within the model thanks to the information provided by short and long-run inflation expectations from Consensus Forecasts. This trend represents the stable inflation path when both the unemployment and the participation gaps are closed. Lastly, we also allow for inflation inertia, so that the Philips curve features both a backward and a forward-looking component like in Galì and Gertler [1999].

Formally, our starting point is the following equation, where the dependent variable is the price inflation gap, that is the difference between realized price inflation and the estimated trend:

$$\pi_t - \pi_t^* = \gamma (\pi_{t-1} - \pi_{t-1}^*) - \gamma \sigma_{\pi^*} \epsilon_t^{\pi^*} - \kappa^u E_t \sum_{T=t}^{\infty} \beta^{T-t} x_T^u + \kappa^p E_t \sum_{T=t}^{\infty} \beta^{T-t} x_T^p - \beta \frac{1 - \rho_a}{1 - \beta \rho_a} \Delta a_t, \quad (8)$$

where γ captures inflation inertia, and κ^u and κ^p denote the reaction of inflation to the present discounted value of future unemployment and participation gaps, respectively.

¹⁷Erceg and Levin [2014] provide a theoretical underpinning for the inclusion of the participation margin in the Phillips curve, introducing a broad definition of employment gap that could be rewritten as follows: $e - e^* = (1 - u)(p - p^*) - p(u - u^*)$.

Notice that we expect inflation to depend negatively on unemployment gaps and positively on participation gaps. The shock $\epsilon_t^{\pi^*}$ indicates a revision of the inflation trend. For instance the inflation trend may shift downward because monetary policy steers long-term inflation expectation towards lower values. Once the shock occurs, it may take some time (depending on the inertia parameter γ) for actual inflation to adjust, so that the inflation gap widens (holding other factors constant).

To make progress on the estimation of equation (8) we need to make some parametric assumptions. We assume that the inflation trend follows a random walk and the unemployment and participation gaps are represented by AR(2) processes. Moreover, the rational expectations hypothesis implies that short and long-term inflation expectations are consistent with the forward iteration of eq. (8) with a margin of error.

We can estimate the unemployment and the participation gaps recalling the decomposition of the realized unemployment and participation rates introduced in Section 2:

$$u_t = x_t^u + z_t^u + \bar{u}_t \quad (9)$$

$$p_t = x_t^p + z_t^p + \bar{p}_t \quad (10)$$

where $z_t^u = u_t^* - \bar{u}_t$ is the deviation of the natural unemployment rate from the structural unemployment rate and $z_t^p = p_t^* - \bar{p}_t$ is the deviation of the natural participation rate from the structural participation rate. We assume that both z_t^u and z_t^p follow an AR(1) process:

$$z_t^u = \rho_{z^u} z_{t-1}^u + \sigma_{z^u, \varsigma} \sigma_{\varsigma} \epsilon_t^{z^u} \quad (11)$$

$$z_t^p = \rho_{z^p} z_{t-1}^p + \sigma_{z^p, \varsigma} \sigma_{\varsigma} \epsilon_t^{z^p}. \quad (12)$$

Equations (11-12) imply that the u^* and p^* converge to their respective structural rates \bar{u}_t and \bar{p}_t in the long-run; however, in the short-run, deviations are allowed with degrees of persistence ρ_{z^u} and ρ_{z^p} . Since σ_{ς} represents the volatility of inflation due to supply shocks,¹⁸ $\sigma_{z^u, \varsigma}$ and $\sigma_{z^p, \varsigma}$ can be interpreted as the signal-to-noise ratios, that is the volatility of the unobserved states u_t^* and p_t^* relative to inflation. Notice that, from the perspective of the Phillips curve model, the trend unemployment and participation rates are exogenous observable inputs; for this reason, the shocks moving z_t^u and z_t^p are fully reflected in the natural unemployment and participation rates. Equations (8) to (12) together with the others described in Appendix B allow us to jointly estimate the parameters of the Phillips curve, u^* and p^* .¹⁹ More specifically, the observable variables are inflation, the unemployment rate, the participation rate, and inflation expectations. The unobserved state variables, which are estimated together with the model parameters

¹⁸If we define $\varsigma_t = -\beta \frac{1-\rho_a}{1-\beta\rho_a} \Delta a_t = \rho_a \varsigma_{t-1} + \sigma_{\varsigma} \epsilon_t^{\varsigma}$ inflation is affected by shocks of volatility σ_{ς} .

¹⁹In our baseline augmented specification the unemployment and the participation gaps are assumed to be uncorrelated. In a robustness exercise we allow them to be correlated through the variance-covariance matrix of their error terms. The estimated correlation is slightly positive but the other results are little affected by this change (Figure C.16 in the Appendix).

through the Kalman filter, include the natural unemployment and participation rates, the inflation trend and the proxy for supply-type inflation pressures.

In order to use all the available information, our baseline model further includes three wage measures: wage per hour in the private sector, wage per equivalent unit of labor in the private sector and negotiated wages. Following Crump et al. [2019], we assume that wage and price inflation are tied by the following relationship: $\pi_t^w = \pi_t + \Delta a_t$. We further assume that real wages grow at rate g_w . We thus add to the model three measurement equations, one for each wage variable:

$$\pi_t^{w^j} = \delta_j (g_w + \pi_t + \Delta a_t) + oe_t^{w^j} \quad \text{with} \quad j = 1, \dots, 3$$

where $\pi_t^{w^j}$ denotes the growth rate of the j -th nominal wage measure, g_w is the constant mean growth rate of real wages and $oe_t^{w^j}$ is an i.i.d. normally distributed measurement error. Like in Crump et al. [2019], we assume that wages and prices react in the same way to the the unemployment and the participation gaps, thus helping to identify the related Phillips curve coefficients. Each wage measure is linked to the others through the scale factor δ_j : $\pi_t^{w^j} = \delta_j \pi_t^{w^1}$, with δ_1 normalized to 1.

We estimate the model with Bayesian techniques using Italian data over the period 1996Q1-2018Q4.²⁰ We use inflation expectations 4 quarters ahead and from 5 to 10 years ahead.²¹

4.2 Results

In what follows we report results based on the structural unemployment and participation rates presented in Section 3, estimated with Kalman filter techniques. The estimation results confirm that the model is well specified. First, the data contain sufficient information to identify the parameters of interest, as their posterior distribution is not fully determined by the prior (Figures C.9-C.10 in the Appendix). Second, the shocks of the main variables included in the model are approximately uncorrelated, homoscedastic and Gaussian (Figures C.11-C.12 in the Appendix).

The priors and the posterior estimates of the model parameters are described in Table 1 and compared to those obtained from an analogous Phillips curve including only the unemployment gap. The inflation process displays a moderate degree of inertia (the median estimate of γ is 0.28), very close to Cogley and Sbordone [2008] estimates on post-Volcker US data when they account for a time-varying inflation trend. The deviations of u^* and p^* from their respective long-run trends are highly auto-correlated, as ρ_{zu} and ρ_{zp} are very close to 1. The median estimate of the signal-to-noise ratio is substantially higher for p^* (the median estimate for $\sigma_{z^p, \varsigma}$ is 0.43) than for the u^* (the median estimate of $\sigma_{z^u, \varsigma}$ is 0.14) implying larger deviations of p^* from its long-run trend.

Let us now turn to the most interesting parameters, those capturing the reaction of prices and wages to the unemployment and participation gaps. κ^u , the coefficient on the

²⁰The choice of the sample period is motivated by data availability, since national accounts are released from 1995 onwards.

²¹Consensus Forecast is available from 1989 onwards.

discounted sum of future unemployment gaps, is relatively small (the median is 0.005); however, the implied overall reaction to the current and lagged unemployment gap (K^u) is higher (the median is 0.084), not far from the estimates of Eser et al. [2020] for the euro area.²² The median estimates of κ^u and K^u in the Phillips curve including only the unemployment gap are higher than those obtained in our augmented model. The estimated median reaction of inflation to the participation gap (κ^p and K^p) is stronger than the impact of the unemployment gap. To understand which margin is more relevant for price dynamics, besides looking at estimated coefficients, we can analyse the historical decomposition of the estimated inflation gap (the dependent variable of the Phillips curve). Figure 7 shows that the participation gap (purple bars) accounts for 70% of the contribution provided by total slack (sum of yellow and purple bars), becoming especially relevant in the low-inflation environment ensuing the double-dip recession: the participation margin alone accounted for roughly one third of the negative inflation gap from 2013 onward.

Figures 8 and 9 plot the natural unemployment and participation rates estimated through the baseline Phillips curve regression including both the unemployment and the participation gap.²³ While u^* is very smooth and well anchored to the trend unemployment rate, p^* follows more closely the observed participation rate.²⁴ As a result, the unemployment gap is more volatile than the participation gap (Figure 10, where the participation gap is reported on a reverse scale so that positive values signal a slack labor market, likewise the unemployment gap). Up to 2013 the estimated participation gap slightly mitigates the indications coming from the unemployment gap. From that period onward, instead, the two margins reinforce each other, as both the unemployment and the participation gaps signal a substantial degree of slack. In the last six years of the sample period the participation margin accounted for almost 30% of total slack but its contribution to price dynamics, as we saw before, was much larger (around 70% of the contribution provided by labor market slack) due to the higher estimated coefficient in the Phillips curve. That in recent years the participation margin has signaled additional slack compared to traditional measures may look surprising given the fast increase in the actual participation rate over the same period (almost 2 p.p., more than what realized in the previous 10 years). However, as Figure 9 makes clear, the rise in participation reflected only partially the strong increase in \bar{p} and p^* ; therefore, the participation gap widened. Further inspection reveals that the recent rise in participation was driven by

²² K^u and K^p denote the overall reaction of inflation to current and lagged unemployment and participation gaps, respectively: $K^u = \kappa^u(\omega_{\pi,1}^u + \omega_{\pi,2}^u)$, $K^p = \kappa^p(\omega_{\pi,1}^p + \omega_{\pi,2}^p)$ (see equation (B.8) in the Appendix). Notice that many other Phillips curve models (like that used by Eser et al. [2020]) do not feature our rich dynamic structure and inflation depends only on the contemporaneous or lagged measure of slack. Therefore that single coefficient should be compared to our composite coefficient K^u .

²³Appendix Figure C.13 further shows that the inflation trend is precisely estimated and follows closely the long-term inflation expectations. Although the model is not designed to capture a specific supply-side factor affecting price dynamics, we find a moderately positive correlation between the estimated Δa_t and productivity growth, consistent with its interpretation (Figure C.14 in the Appendix).

²⁴One possible explanation is related to the assumption that it takes time for the natural rates to converge to their structural values: while the structural unemployment rate has remained constant in the last decades, the structural participation rate has increased substantially starting from the end of the nineties. It may therefore still take time for p^* to fully incorporate such changes, while u^* is already around the relatively constant structural unemployment rate.

the reduction in the exit rate of the elderly induced by a pension reform enacted in 2012 (see next Section). The participation rates of the youngsters and prime-age men, however, hovered around cyclically low levels (Figure C.8), thus pointing to more slack than what suggested by aggregate figures. Our framework is able to detect these changes in a parsimonious way through the estimation of the natural participation rate.

The results of our baseline augmented Phillips curve can be compared to those derived from alternative, simpler models. First of all, we consider a Phillips curve regression where inflation positively depends on the employment gap, defined as the difference between the observed employment rate e_t and the unobserved estimated value e_t^* . Second, we consider the standard model including only the unemployment gap. On the one hand, if the participation margin has no explanatory power on inflation dynamics we expect the three specifications to yield similar results and our baseline model to underperform in terms of estimation precision due to its greater complexity. On the other hand, if the participation margin is relevant but its effect is akin to unemployment, the employment gap would summarize all the relevant information. To compare these different models we combine the unemployment and participation gaps of our baseline setup to obtain the corresponding employment gap²⁵: $e - e^* = (1 - u)(p - p^*) - p(u - u^*)$. The results are shown in Figure 11, where the measures of employment gap are divided by the actual participation rate to make them comparable with the unemployment gap. The gap estimated under our baseline augmented model displays some notable differences with the other two. Focusing on the most recent period, the combination of unemployment and participation gap signals considerably more slack compared to the estimates based solely on unemployment (the peak level recorded in 2015 is about four times larger) but slightly less if compared to the model based on employment.²⁶

Notice that these results do not necessarily imply that the standard Phillips curve does not detect any significant pressure of labor market slack on inflation. Indeed, also the standard model points to a significant negative contribution of the unemployment gap to price dynamics in recent years because the estimated Phillips curve coefficient is higher (see Appendix Figure C.15). Therefore, we further assess the relative performance of the augmented model by asking which estimated measure of slack can better explain observed price dynamics from an *ex-ante* perspective. In the spirit of Banbura et al. [2015], we compute inflation forecasts from 2014q4 onward conditional on our estimated measures of slack. We find that our baseline augmented model which explicitly takes into account the participation gap outperforms the specification based either on the combined employment gap or solely on the unemployment gap (Figure 12). Indeed, the negative participation gap crucially adds to overall slack, putting downward pressures on inflation.

Finally, in the same vein as Crump et al. [2019] and Davins and Primiceri [2019] we also investigate the distinct roles played by inflation expectations and labor market flows. We find that omitting the information on inflation expectations increases the uncertainty around our estimates but has only a moderate impact on the median evolution

²⁵See Erceg and Levin [2014] for a similar decomposition.

²⁶Appendix Figure C.16 shows the results for the model that allows unemployment and participation gaps to be correlated.

of economic slack. The estimates of u^* and p^* , instead, change substantially when they are not anchored to the respective structural rates extracted from labor market flows (Figure C.17 in the Appendix).²⁷

5 The effects of an unexpected pension reform

In this section we consider a pension reform that raised statutory retirement age and exogenously shifted labor supply of the elderly without directly affecting unemployment or job search intensity. We believe that this is an ideal setting to fully exploit our estimation method. Our approach indeed is perfectly suited to estimate whether this reform had implications for labor market slack and to quantify them. Moreover, by estimating the impact of the reform on p^* and u^* we can compute its overall effect on potential employment, i.e., the level of employment relevant for potential output.

In particular, we use our framework to evaluate the impact on potential employment of the Fornero reform, a far reaching and unexpected pension reform that took place in Italy at the height of the sovereign debt crisis (in 2012), motivated by the need of limiting pension expenditures.²⁸ It consisted of a slight reduction in the pension transfer and a substantial increase in the statutory retirement age (the maximum delay was equal to seven years for certain workers' categories). In our context, such a reform would increase the structural participation rate through a change in the exit margin, due to both the decrease in the outside option (pension income) and the increase in the minimum statutory age for retirement. Our goal is to quantify the macro effect of this reform on natural participation and unemployment rates, and to assess its implications for price dynamics. The richness of our approach, that relies on detailed micro-level data and that adds a third labor market state –inactivity– provides the ideal setting to estimate these effects. In the research design, we leverage the fact that the reform was unexpected and implemented a few days after its announcement; as a consequence its effects were magnified, as firms' turnover and workers' labor supply decisions were not affected in the years preceding its implementation. We use this element in the evaluation of the impact of the reform, by deriving our best forecast of natural and structural rates for the 2011-2015 period based on the pre-reform trends in labor market flows. In order to do that, we first repeat the estimation outlined in Sections 3 and 4 using only pre-reform data (until 2011) and we project structural and natural rates until 2015.²⁹ Concerning the structural rates, we project them assuming that in the absence of the reform each group specific rate would have behaved similarly to what observed in the years before the reform. We denote as $\{\bar{u}_t^{PRE}\}_{t=2012}^{2015}$ and $\{\bar{p}_t^{PRE}\}_{t=2012}^{2015}$ the paths of the structural unemployment and structural participation in place prior to the reform; we interpret

²⁷In this alternative model the structural unemployment and participation rates are estimated within the same Phillips curve framework assuming a random-walk process.

²⁸For an overview of the reform see [Carta et al., 2021] and [Carta and De Philippis, 2021].

²⁹The reform was passed in December 2011 and started to have effect in January 2012. We assume that the reform produced (most of) its effects by 2015, as the average delay in retirement age for 55+ workers was less than 3 years (Carta et al. [2021]). Enlarging the time window may be problematic as we would likely be capturing other shocks not related to the reform.

them as the counterfactual evolution of the Italian labor market throughout the period 2012–2015 absent the reform. Then, we estimate again the structural rates using the data until 2015, i.e. including the post reform period, and we obtain $\{\bar{u}_t^{POST}\}_{t=2012}^{2015}$ and $\{\bar{p}_t^{POST}\}_{t=2012}^{2015}$; we normalize their levels in 2011 to be equal to those estimated in the pre-reform period.³⁰ Finally, we compute the effect of the reform as the simple difference between the POST and the PRE-series.

According to our estimates, the reform was successful in increasing the structural participation of the older groups, by about 3 p.p. and 1 p.p. for males and females aged 55–64 respectively, and by 1 p.p. for females aged 35–54 (see Figure C.18). Moreover, we find that the impact of the reform on structural unemployment was negligible (Figure C.19). Overall, these effects imply that the policy increased the aggregate structural participation rate by about 0.7 p.p. in 2015, while the trend unemployment rate remained roughly unchanged, as shown in Figures 13 and 14. Reflecting the differential impact across groups, the total effect on aggregate participation was mainly due to the 55–64 groups (that explain about two thirds of the overall variation), with almost all the rest accounted for by females aged 35–54 (see Figure C.20).³¹ We conclude that the Fornero reform had a substantial impact on structural participation, with a negligible effect on the structural unemployment rate.

Finally, we further investigate whether the Phillips curve model detects any change in the level of the natural unemployment and participation rates before and after the implementation of the pension reform. In order to do that, we estimate the model on the pre-reform period (1996Q1–2011Q4) and obtain estimates for the parameters and the evolution of the state variables based exclusively on pre-reform data. Then we project the model forward using these estimated laws of motion and assuming that the only data that become available in the post-reform period are those related to the structural unemployment and participation rates.³² By feeding the model with the series of $\{\bar{u}_t^{PRE}\}_{t=2012}^{2015}$ and $\{\bar{p}_t^{PRE}\}_{t=2012}^{2015}$ we obtain the evolution of the natural rates absent the reform, while feeding the model with $\{\bar{u}_t^{POST}\}_{t=2012}^{2015}$ and $\{\bar{p}_t^{POST}\}_{t=2012}^{2015}$ delivers the projections of the natural rates with the pension reform in place. In line with our results on the structural rates, we find negligible effects of the pension reform on u^* (Figure 15) and more substantial effects on p^* (Figure 16). According to our median estimates, in 2015 the natural participation rate would have been 0.7 p.p. lower absent the Fornero reform. This implies, everything else equal, that a given observed participation rate is now associated to lower inflation pressures due to the reform. This explains why the relatively strong increase in the participation rate since 2012 did not have a positive impact on wage growth: the even larger rise in p^* driven by the pension reform resulted in a negative participation gap, thus contributing to subdued inflation. Overall, our estimates indicate that the reform

³⁰As we want to avoid that our filtering technique changes the estimation for the period prior to the reform, for each subgroup we define an adjustment factor $adj_g^x = \bar{x}_{g,2011q4}^{PRE} - \bar{x}_{g,2011q4}^{POST*}$, and construct our final POST-series accordingly: $\bar{x}_{g,t}^{POST} = \bar{x}_{g,t}^{POST*} + adj_g^x$ for $x \in \{u, p\}$ and $t \in [2012, 2015]$, where $\bar{x}_{g,t}^{POST*}$ refers to the non-adjusted series.

³¹This result is consistent with Carta and De Philippis [2021], who look at whether the reform affected also labor supply of middle-aged individuals and find a positive and significant effect.

³²Notice that we do not use any data from the post-reform period, most likely contaminated by other shocks, except from the structural rates that have already been purged of their cyclical component.

increased potential employment by 1.1%.

6 Conclusions

Motivated by the magnitude and cyclicity of the transitions into and out of the labor force, we investigate the role of the participation margin for the measurement of labor market slack and for price dynamics. Extending Crump et al. [2019]’s work, we propose a unified framework for the estimation not only of the natural unemployment, but also of the natural participation rate, based on a forward-looking Phillips curve informed by labor market flows between employment, unemployment and inactivity among different demographic groups.

We focus on Italy for two main reasons. First, fluctuations into and out of the labor force are much larger than those between employment and unemployment, also in comparison to other advanced countries. Second, a far-reaching and unexpected rise in statutory retirement age taking place in 2012 dramatically increased elderly workers’ labor supply by decreasing their flows from employment into inactivity; this provides an ideal setting to tease out the implications for the measurement of labor market slack of an exogenous shift in labor supply not directly related to unemployment or job search intensity.

Our results show that the inclusion of the participation margin is crucial for the measurement of overall labor market slack and helps explaining the low inflation regime that followed the Sovereign Debt Crisis, both when performing an historical decomposition and when conducting a conditional forecast exercise.

By analyzing the effects of the pension reform, we find that the large increase in participation of older individuals, determined by the change in the full retirement age, did not generate any change in structural activity for younger age classes or in structural unemployment. Overall, the natural unemployment rate was unaffected, while the natural participation rate increased by 0.7 percentage points. Failing to account for changes in natural participation would have thus led to underestimate the degree of economic slack after the reform and produced a positive bias in inflation forecasts.

Taken together, these results show that the participation margin matters for price dynamics in Italy. Using Eurostat Labor Force Survey microdata for the period 1992-2019, we find that yearly transitions into and out of the labor force by gender and age show similar patterns in the five largest European countries (France, Germany, Italy, Spain, United Kingdom)³³, while their incidence greatly varies, also over time.³⁴ Checking whether such patterns have implication for price and wage dynamics is certainly a promising avenue for future research.

³³In particular, we regress for each country the probability of entering (exiting) the labor force at one year intervals on gender, three age classes and year dummies. Albeit coefficient point estimates are statistically different, they depict similar regularities across countries: Men are more likely to both enter or exit the labor force, while - as expected - entry is more (less) likely than exit for young (elderly) workers.

³⁴All results are reported in figures XX of the Appendix.

References

- Daniel Aaronson, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan. Declining Labor Force Participation and its Implications for Unemployment and Employment Growth. *Chicago FED Economic Perspectives*, 38(4), 2014.
- Katharine G. Abraham, John C. Haltiwanger, and Lea E. Rendell. How Tight is the U.S. Labor Market? *Brookings Papers on Economic Activity*, 2020.
- Daron Acemoglu and Pascual Restrepo. Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation. *American Economic Review*, 107(5): 174–79, May 2017.
- Laurence Ball and Sandeep Mazumder. Inflation Dynamics and the Great Recession. *Brookings Papers on Economic Activity*, 42(1):337–405, 2011.
- Laurence Ball and Sandeep Mazumder. A Phillips Curve with Anchored Expectations and Short-Term Unemployment. *Journal of Money, Credit and Banking*, 51(1):111–137, 2019.
- M. Banbura, D. Giannone, and M. Lenza. Conditional Forecasts and Scenario Analysis with Vector Autoregressions for Large Cross-sections. *International Journal of Forecasting*, 2015.
- Regis Barnichon and Andrew Figura. Labor Market Heterogeneity and the Aggregate Matching Function. *American Economic Journal: Macroeconomics*, 7(4):222–49, 2015.
- Regis Barnichon and Geert Mesters. On the Demographic Adjustment of Unemployment. *Review of Economics and Statistics*, 100(2):219–231, 2018.
- Regis Barnichon and Geert Mesters. Identifying Modern Macro Equations with Old Shocks. *The Quarterly Journal of Economics*, 135(4):2255–2298, 2020.
- David N.F. Bell and David G. Blanchflower. Underemployment in the UK Revisited. *National Institute Economic Review*, 224:F8–F22, 2013.
- Nicola Bianchi, Giulia Bovini, Jin Li, Matteo Paradisi, and Michael L. Powell. Career Spillovers in Internal Labor Markets. Review of economic studies, National Bureau of Economic Research, Inc, March 2022.
- Olivier Blanchard and Diamond. The Cyclical Behavior of the Gross Flows of U.S. Workers. *Brookings Papers on Economic Activity*, 21(2):85–156, 1990.
- Elena Bobeica and Marek Jarociński. Missing Disinflation and Missing Inflation: A VAR Perspective. *International Journal of Central Banking*, 15(1):199–232, March 2019.
- Isabel Cairó, Shigeru Fujita, and Camilo Morales-Jimenez. The Cyclical Flow of Labor Force Participation: The Role of Labor Supply Elasticities and Wage Rigidity. *Review of Economic Dynamics*, 43:197–216, 2022.

- Alessia Campolmi and Stefano Gnocchi. Labor Market Participation, Unemployment and Monetary Policy. *Journal of Monetary Economics*, 79(C):17–29, 2016.
- Francesca Carta and Marta De Philippis. Working Horizon and Labour Supply: the Effect of Raising Retirement Age on Middle-aged Individuals. Temi di Discussione 1314, Bank of Italy, 2021.
- Francesca Carta, Francesco D’Amuri, and Till von Wächter. Workforce Aging, Pension Reforms, and Firm Outcomes. *NBER Working Paper*, 2021.
- Sekyu Choi, Alexandre Janiak, and Benjamín Villena-Roldán. Unemployment, Participation and Worker Flows over the Life-cycle. *The Economic Journal*, 125(589):1705–1733, 2015.
- Timothy Cogley and Argia M. Sbordone. Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98(5):2101–26, December 2008.
- Olivier Coibion and Yuriy Gorodnichenko. Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232, January 2015.
- Richard Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. A Unified Approach to Measuring u^* . *Brookings Papers on Economic Activity*, 2019.
- Richard Crump, Christopher J. Nekarda, and Nicolas Petrosky-Nadeau. Unemployment Rate Benchmarks. Finance and Economics Discussion Series 2020-072, FED, August 2020.
- Francesco D’Amuri, Marta De Philippis, Elisa Guglielminetti, and Salvatore Lo Bello. Slack and Prices during Covid-19: Accounting for Labor Market Participation. *Labour Economics*, 75, 2022.
- Steven J. Davins and Giorgio E. Primiceri. Comments and discussion. *Brookings Papers on Economic Activity*, pages 215–238, 2019.
- Marco Del Negro, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti. What’s Up with the Phillips Curve? *Brookings Papers on Economic Activity*, pages 301–357, 2020.
- Michael Elsby, Bart Hobijn, Fatih Karahan, Gizem Koşar, and Ayşegül Şahin. Flow Origins of Labor Force Participation Fluctuations. In *AEA Papers and Proceedings*, volume 109, pages 461–64, 2019.
- Michael WL Elsby, Bart Hobijn, and Ayşegül Şahin. On the Importance of the Participation Margin for Labor Market Fluctuations. *Journal of Monetary Economics*, 72: 64–82, 2015.

- Niklas Engbom. Firm and Worker Dynamics in an Aging Labor Market. *Minneapolis FED WP*, (756), 2019.
- Christopher J. Erceg and Andrew T. Levin. Labor Force Participation and Monetary Policy in the Wake of the Great Recession. *Journal of Money, Credit and Banking*, 46 (S2):3–49, 2014.
- Fabian Eser, Peter Karadi, Philip R. Lane, Laura Moretti, and Chiara Osbat. The Phillips Curve at the ECB. Working Paper Series 2400, European Central Bank, May 2020.
- Eurostat. Labour market transitions - Quarterly data. Technical report, 2020.
- R. Jason Faberman, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa. The Shadow Margins of Labor Market Slack. *Journal of Money, Credit and Banking*, 2020.
- James Feyrer. Demographics and Productivity. *The Review of Economics and Statistics*, 89(1):100–109, 2007.
- Shigeru Fujita and Garey Ramey. The Cyclicity of Separation and Job Finding Rates. *International Economic Review*, 50(2):415–430, 2009.
- Jordi Galí. The Return of the Wage Phillips Curve. *Journal of the European Economic Association*, 9(3):436–461, 2011.
- Jordi Galí and Mark Gertler. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44(2):195–222, 1999.
- Jordi Galí, Frank Smets, and Rafael Wouters. Unemployment in an Estimated New Keynesian Model. In *NBER Macroeconomics Annual 2011, Volume 26*, NBER Chapters, pages 329–360. National Bureau of Economic Research, Inc, July 2011.
- Pietro Garibaldi and Etienne Wasmer. Equilibrium Search Unemployment, Endogenous Participation, and Labor Market Flows. *Journal of the European Economic Association*, 3(4):851–882, 2005.
- Pedro Gomes. Labour Market Flows: Facts from the United Kingdom. *Labour Economics*, 19(2):165–175, 2012.
- Robert J Gordon. The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery. Working Paper 19390, National Bureau of Economic Research, August 2013.
- Robert E. Hall and Sam Schulhofer-Wohl. Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers. *American Economic Journal: Macroeconomics*, 10(1):1–32, January 2018.
- Bart Hobijn and Ayşegül Şahin. Maximum Employment and the Participation Cycle. *Paper prepared for the Jackson Hole symposium*, 2021.

- Andreas Hornstein, Marianna Kudlyak, and Fabian Lange. Measuring resource utilization in the labor market. *Economic Quarterly, FRB Richmond*, 2020.
- Thomas B King. Variable Participation and Cyclical Unemployment. *Available at SSRN 1868930*, 2011.
- Per Krusell, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. Gross Worker Flows over the Business Cycle. *American Economic Review*, 107(11):3447–76, November 2017.
- Per Krusell, Toshihiko Mukoyama, Richard Rogerson, and Aysegul Sahin. Gross Worker Flows and Fluctuations in the Aggregate Labor Market. *Review of Economic Dynamics*, 37:S205–S226, 2020.
- Marianna Kudlyak and Felipe F Schwartzman. Accounting for Unemployment in the Great Recession: Nonparticipation Matters. *Federal Reserve Bank of Richmond Working Paper*, (12-04), 2012.
- Etienne Lalé. Understanding Fluctuations in the Ins and Outs of the Labor Force. *Available at SSRN 2534941*, 2013.
- Sophocles Mavroeidis, Mikkel Plagborg-Møller, and James H. Stock. Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature*, 52(1):124–88, March 2014.
- Michael McLeay and Silvana Tenreyro. *Optimal Inflation and the Identification of the Phillips Curve*, pages 199–255. University of Chicago Press, June 2019.
- Dale T Mortensen and Christopher A Pissarides. Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415, 1994.
- Barbara Petrongolo and Christopher A Pissarides. The Ins and Outs of European Unemployment. *American Economic Review*, 98(2):256–62, 2008.
- Michael Pries and Richard Rogerson. Search Frictions and Labor Market Participation. *European Economic Review*, 53(5):568–587, 2009.
- Alfonso Rosolia. The Italian labor market after the Great Crisis. in “Gli Effetti della Crisi sul Potenziale Produttivo,” Vol. 18, Workshops and Conferences, Bank of Italy, 2014.
- Paolo Sestito. Flussi e Mobilità nel Mercato del Lavoro Italiano (1979-1988). *Ricerche Economiche*, 17:143–171, 1988.
- Robert Shimer. Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics*, 15(2):127–148, 2012.
- James H. Stock and Mark W. Watson. Slack and Cyclically Sensitive Inflation. *Journal of Money, Credit and Banking*, 52(S2):393–428, 2020.

K. Strand and T. Dernburg. Cyclical Variation in Civilian Labor Force Participation. *The Review of Economics and Statistics*, 46:378–391, 1964.

Murat Tasci. The Ins and Outs of Unemployment in the Long run: Unemployment Flows and the Natural Rate. Koç University-TÜSIAD Economic Research Forum Working Paper Series, 2012.

Tables and Figures

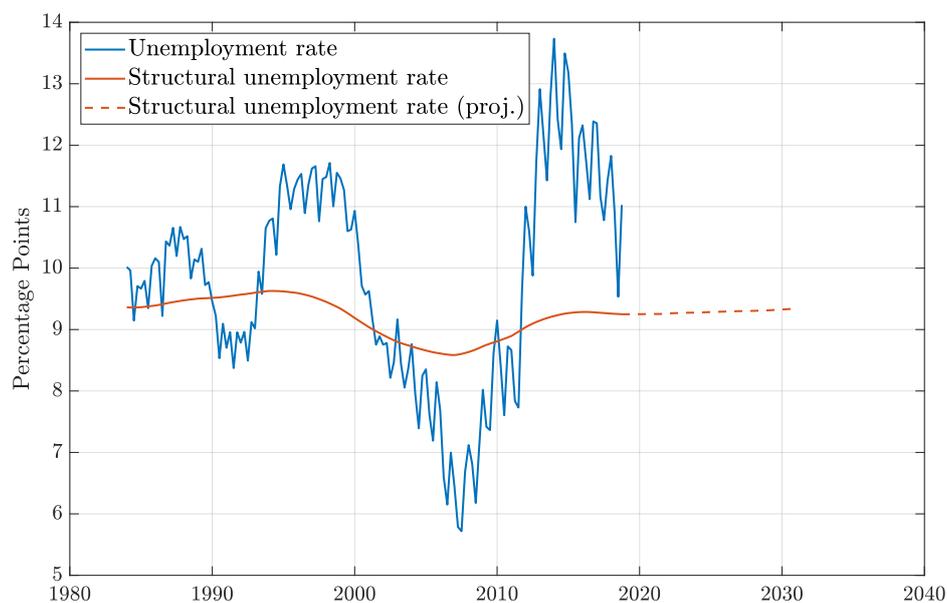


Figure 1: Aggregate Structural Unemployment Rate.

Note: The figure plots the quarterly unemployment rate of the population 15–64 (blue line) and the estimated trend unemployment rate (red line) for the years between 1984q1-2018q4. The dashed line represents the projection of the trend unemployment rate until 2030, based on Eurostat demographic projections and assuming that the long term trends in labor market flows will follow the the same dynamics observed between 2015 and 2018.

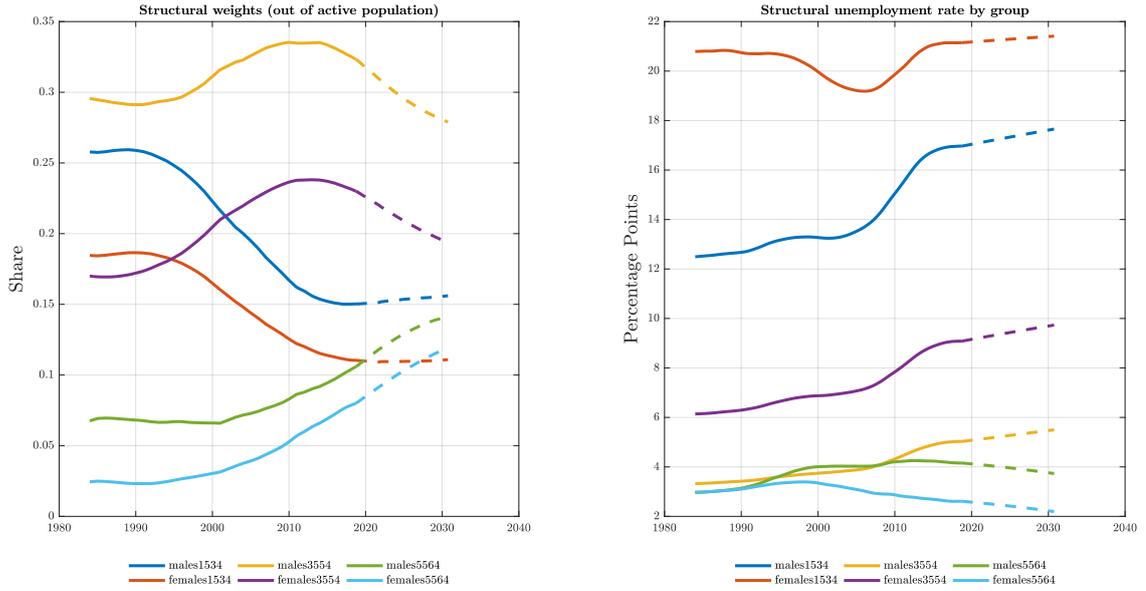


Figure 2: Determinants of Aggregate Structural Unemployment.

Note: The left panel plots the weights $\bar{\omega}_{g,t}^u$ of each demographic group g . These weights are equal to the product between the weight in the population of each group g ($\omega_{g,t}^p$) and the ratio between the group-specific and the aggregate trend participation rate ($\frac{\bar{p}_{g,t}}{\bar{p}_t}$). The right panel displays the trend unemployment rate of each subgroup g . The dashed lines represent projections until 2030, based on Eurostat demographic projections and assuming that the long term trends in labor market flows will follow the same dynamics observed between 2015 and 2018.

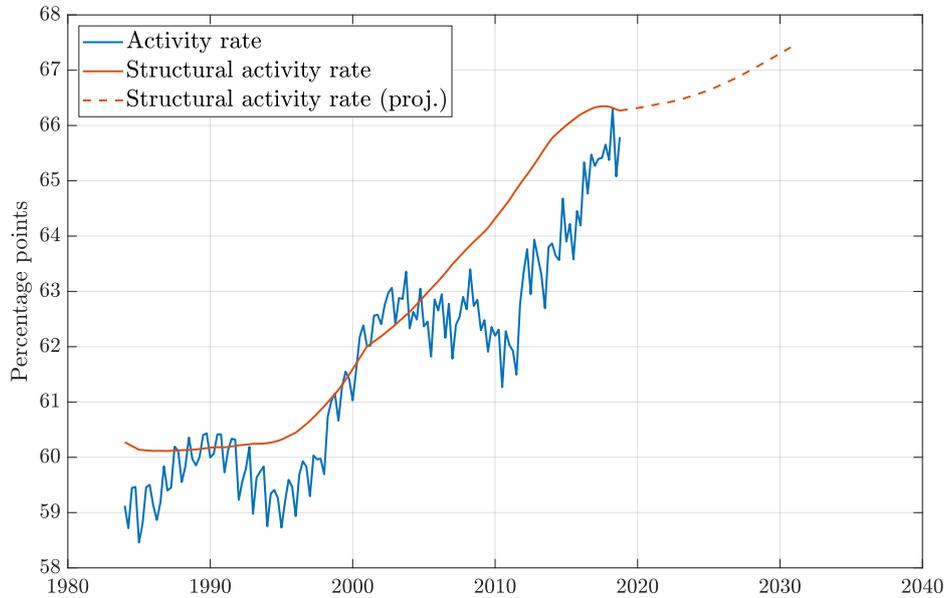


Figure 3: Aggregate Structural Participation Rate.

Note: The figure plots the quarterly activity rate of the population 15–64 (blue line) and the estimated trend activity rate (red line) for the years between 1984q1-2018q4. The dashed line represents the projection of the trend activity rate until 2030, based on Eurostat demographic projections and assuming that the long term trends in labor market flows will follow the same dynamics observed between 2015 and 2018.

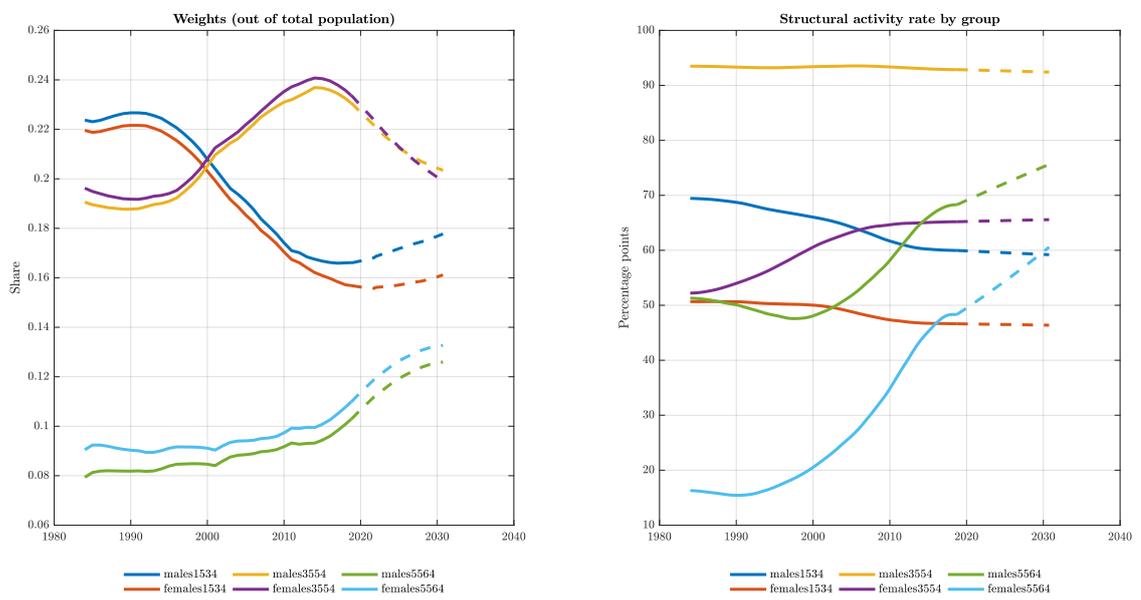


Figure 4: Determinants of Aggregate Structural Participation.

Note: The left panel plots the population weights $\omega_{g,t}^p$ of each demographic group g . The right panel displays the trend activity rate of each subgroup g . The dashed lines represent projections until 2030, based on Eurostat demographic projections and assuming that the long term trends in labor market flows will follow the same dynamics observed between 2015 and 2018.

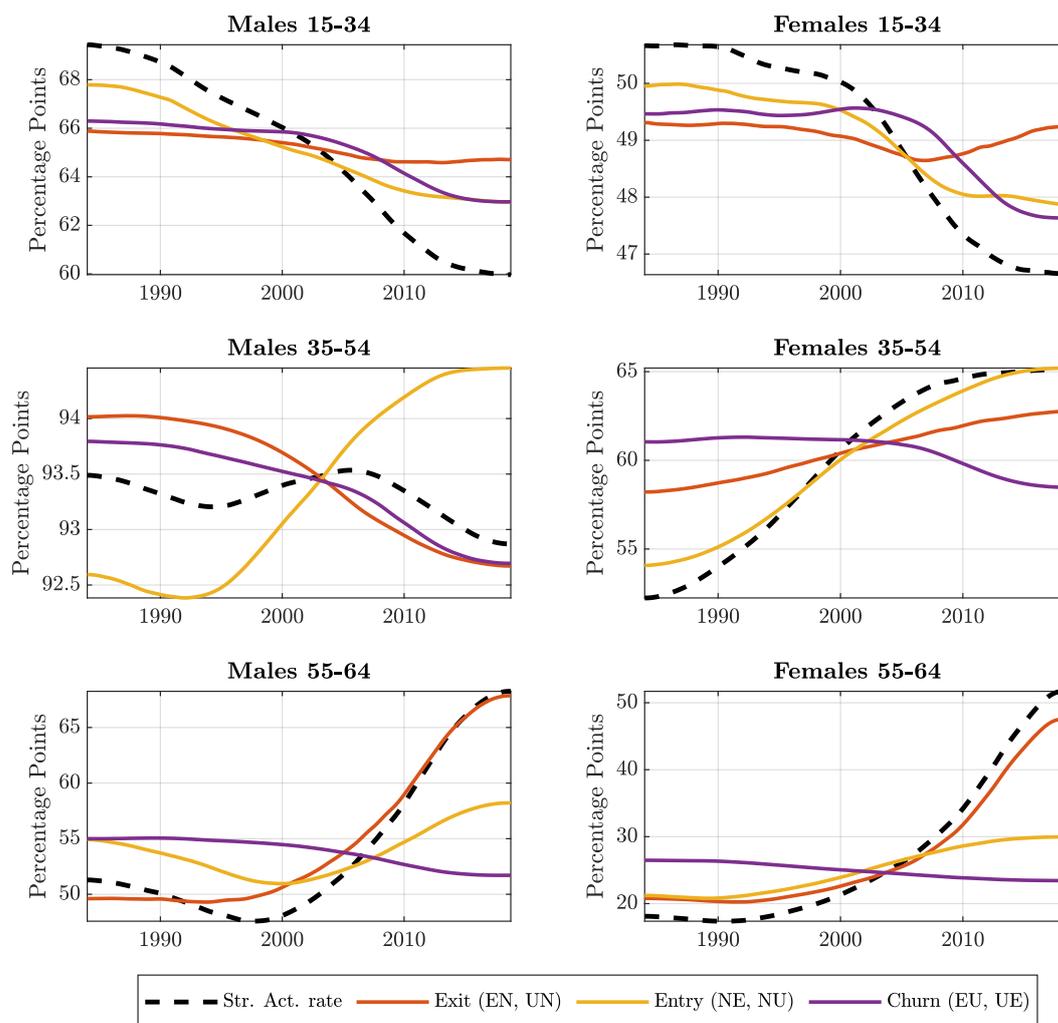


Figure 5: Decomposition of group-specific structural activity rates.

Note: The figure plots the group-specific estimated trend activity rates (blue line) for the years between 1984q1-2018q4, and the counterfactual series in which we let vary only a subset of flows at the time (exit: EN, UN; entry: NE, NU; churn: EU, UE), fixing the other flows at their sample mean.

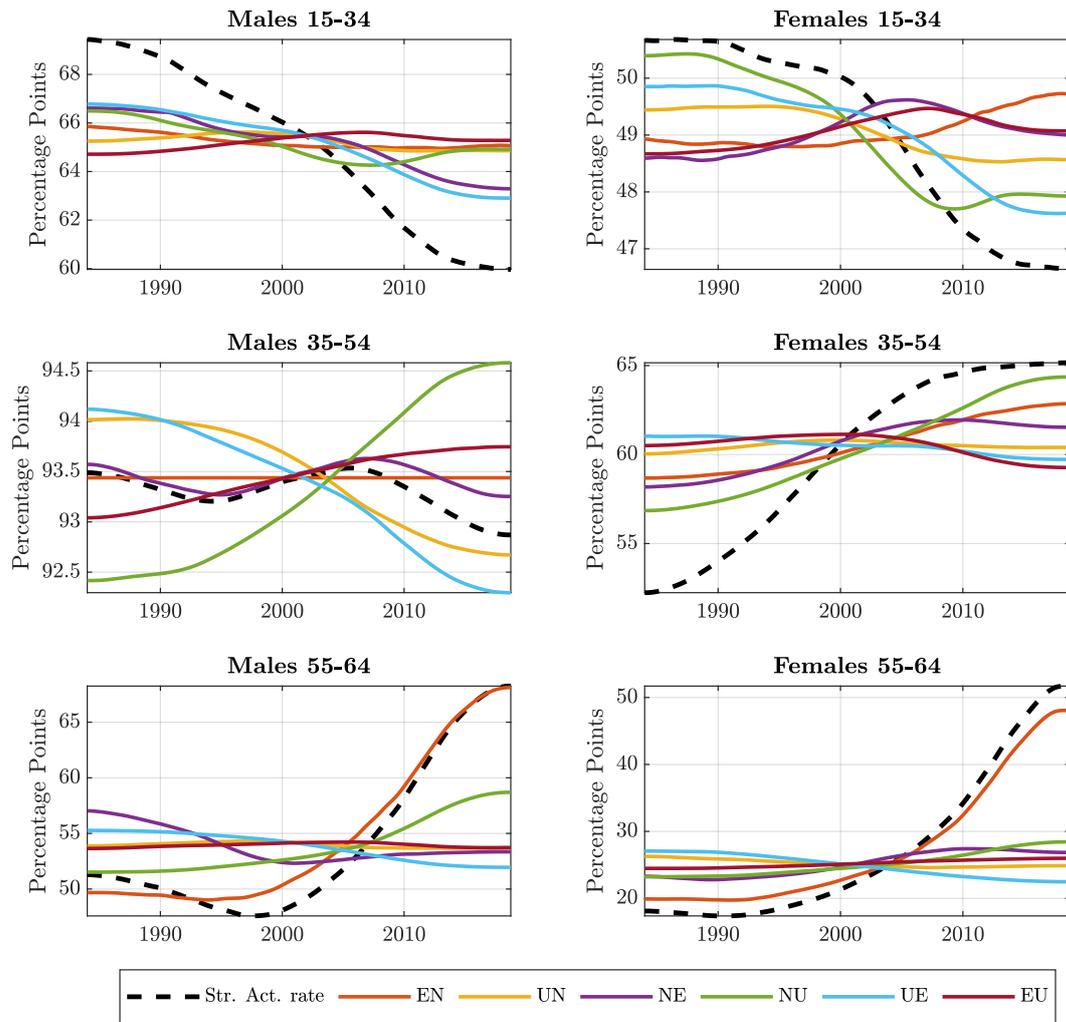


Figure 6: Decomposition of group-specific structural activity rates.

Note: The figure plots the group-specific estimated trend activity rates (blue line) for the years between 1984q1-2018q4, and the counterfactual series in which we let vary only a flow at the time, fixing the other flows at their sample mean.

Table 1: Parameter estimates

	Prior			Augmented PC			Standard PC (UGAP only)		
	Dist.	Mean	Std	Median	5%	95%	Median	5%	95%
$a_{x^u,1}$	Gamma	1.25	0.200	1.11	1.01	1.18	1.04	0.893	1.46
$a_{x^u,2}$	Normal	0.000	1.00	-0.148	-0.219	-0.054	-0.136	-0.529	0.020
κ^u	Normal	0.150	0.100	0.005	0.001	0.012	0.077	0.022	0.221
γ	Beta	0.500	0.265	0.284	0.152	0.375	0.277	0.173	0.395
ρ_{z^u}	Beta	0.950	0.035	0.954	0.897	0.986	0.969	0.942	0.989
ρ_a	Beta	0.500	0.200	0.387	0.299	0.466	0.394	0.220	0.486
$\sigma_{x^u}^2$	InvGamma	0.112	0.000	0.096	0.073	0.126	0.019	0.009	0.040
σ_{ζ}^2	InvGamma	1.00	0.000	0.259	0.200	0.344	0.254	0.196	0.335
$\sigma_{\pi^*}^2$	InvGamma	0.112	0.000	0.014	0.010	0.020	0.014	0.010	0.020
$\sigma_{z^u, \zeta}$	InvGamma	0.150	0.050	0.142	0.085	0.252	0.586	0.497	0.672
g_w	Normal	0.400	0.050	0.384	0.298	0.473	0.370	0.306	0.439
K^u	–	–	–	0.084	0.017	0.167	0.722	0.339	1.43
$a_{x^p,1}$	Gamma	1.25	0.200	0.865	0.798	1.04	–	–	–
$a_{x^p,2}$	Normal	0.000	1.00	0.025	-0.129	0.118	–	–	–
κ^p	Normal	0.150	0.100	0.064	0.016	0.221	–	–	–
ρ_{z^p}	Beta	0.950	0.035	0.966	0.931	0.988	–	–	–
$\sigma_{x^p}^2$	InvGamma	0.112	0.000	0.027	0.013	0.052	–	–	–
$\sigma_{z^p, \zeta}^p$	InvGamma	0.150	0.050	0.428	0.326	0.505	–	–	–
K^p	–	–	–	0.666	0.287	1.30	–	–	–

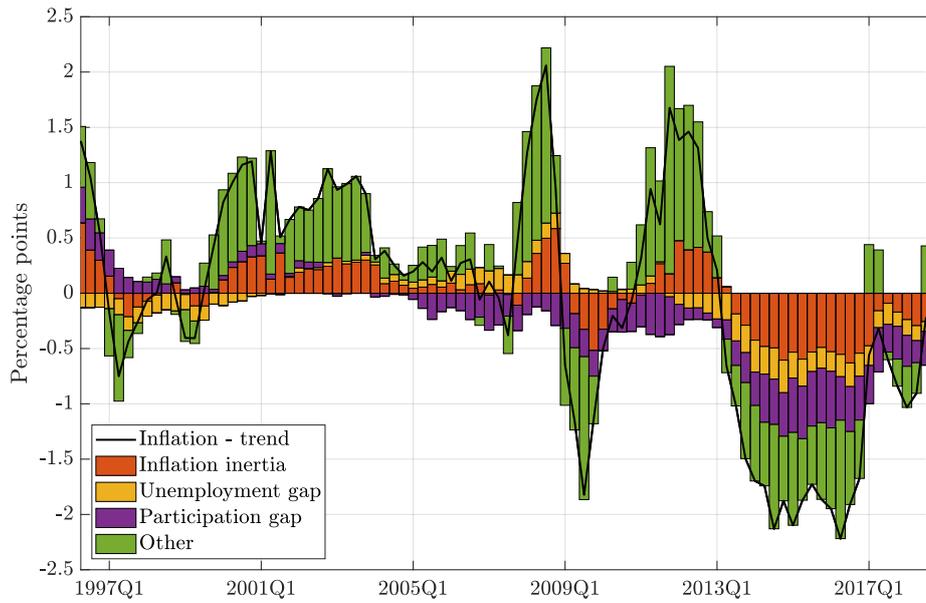


Figure 7: Historical decomposition of the inflation gap in the augmented model

Note: The solid line represents the historical evolution of median inflation gap (realized inflation - estimated trend inflation) and the colored bars the median contributions of the factors included in the augmented Phillips curve model (UGAP + PGAP). The model is estimated over the period 1996Q1–2018Q4.

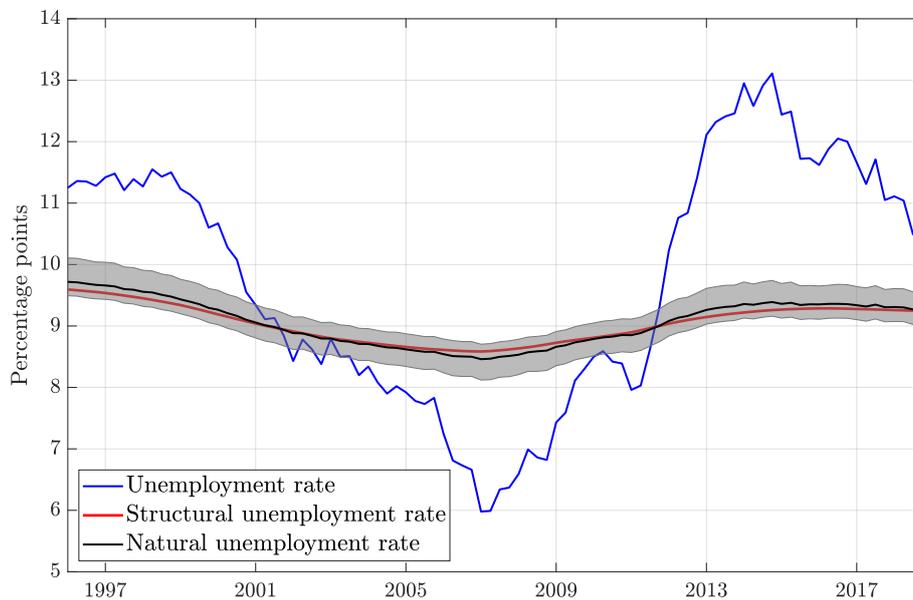


Figure 8: u^* estimated through the augmented Phillips curve model (UGAP + PGAP)

Note: Shading denotes the 68% credible interval. The model is estimated over the period 1996Q1–2018Q4.

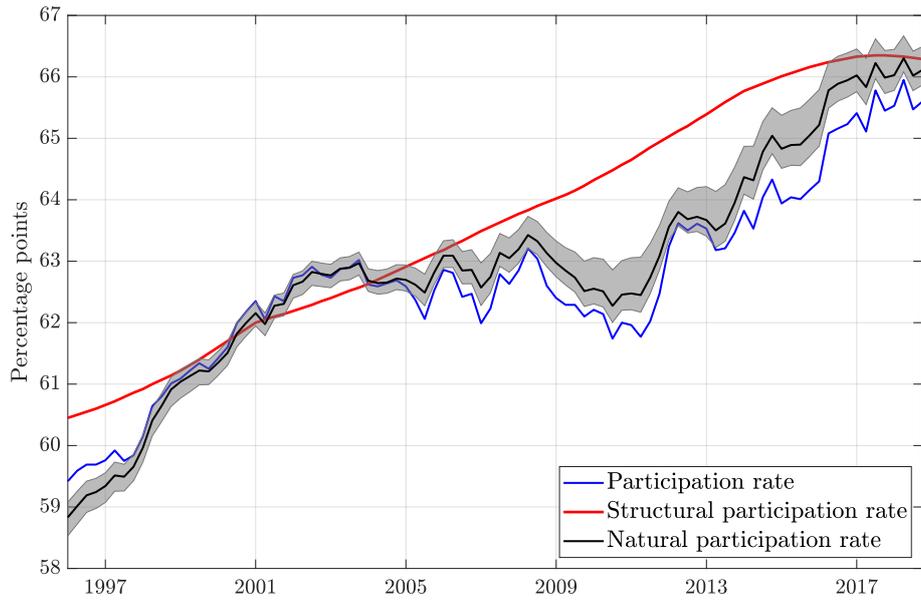


Figure 9: p^* estimated through the augmented Phillips curve model (UGAP + PGAP)

Note: Shading denotes the 68% credible interval. The model is estimated over the period 1996Q1–2018Q4.

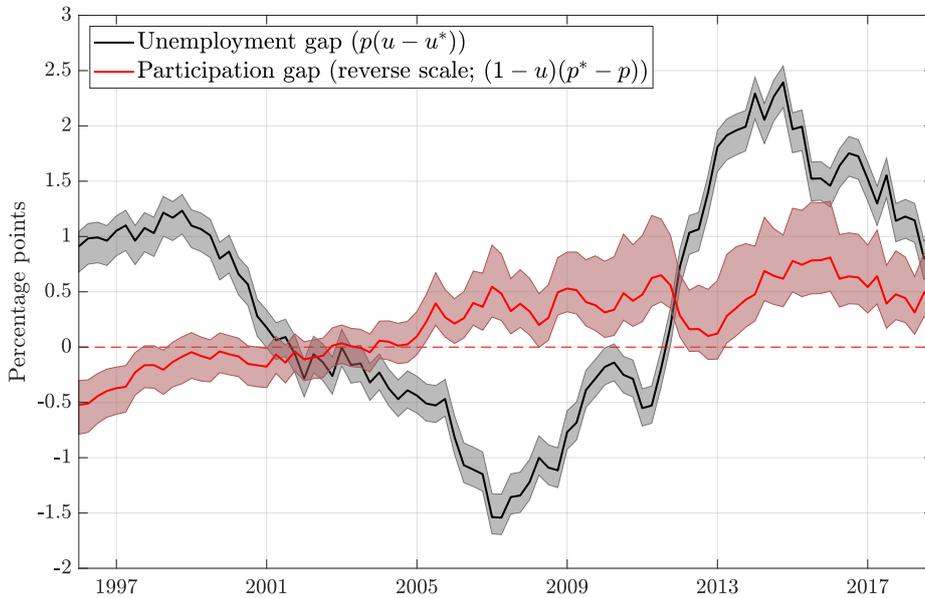


Figure 10: Unemployment and participation gaps estimated through the augmented Phillips curve model

Note: Shading denotes the 68% credible interval. The model is estimated over the period 1996Q1–2018Q4. Positive unemployment and participation gaps denote a slack labor market.

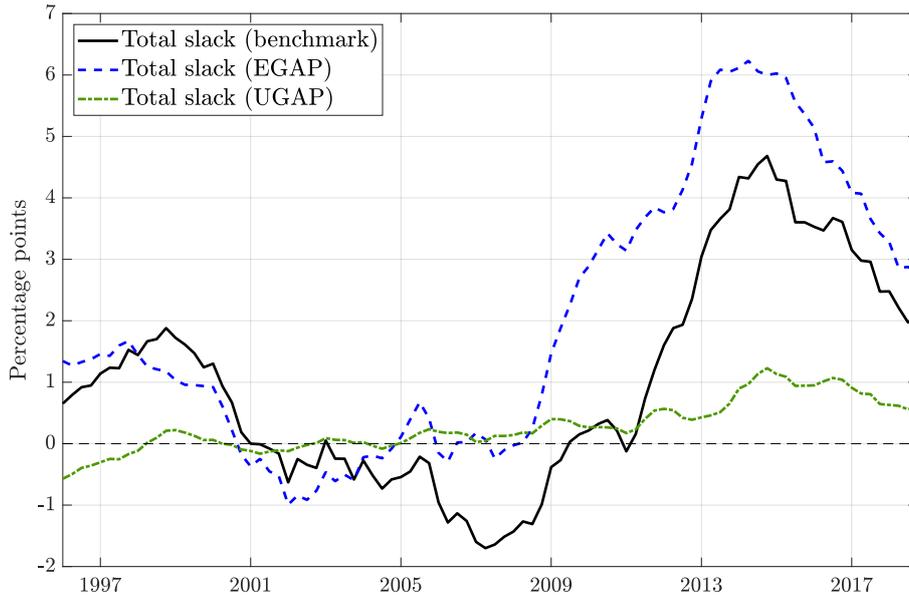


Figure 11: Employment and unemployment gap estimates in different models

Note: The black solid line represents the median employment gap estimated through the baseline augmented Phillips curve including both the unemployment and the participation gap. The blue dashed line refers to the model including only a combined employment gap. The green line represents the median unemployment gap obtained through a Phillips curve including only the unemployment gap. The employment gaps in the first two models are represented on a reverse scale (positive values indicate a slack labor market) and are divided by the actual participation rate to make them comparable with the UGAP derived from the third model.

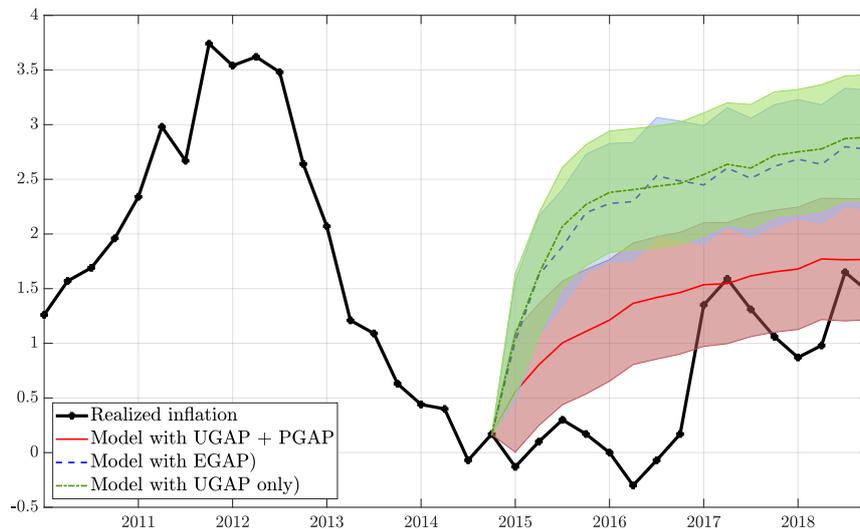


Figure 12: Conditional forecasts of inflation

Note: The red solid red line represents the median conditional forecast of inflation obtained from the baseline Phillips curve including both the unemployment and the participation gap. The blue dashed line refers to the model including only a combined employment gap while the green line refers to the model including only the unemployment gap. The conditioning set includes labor market variables, that is actual, structural and natural unemployment and participation rates from 2015q1 till 2018q4. Shading denotes the 68% credibility interval.

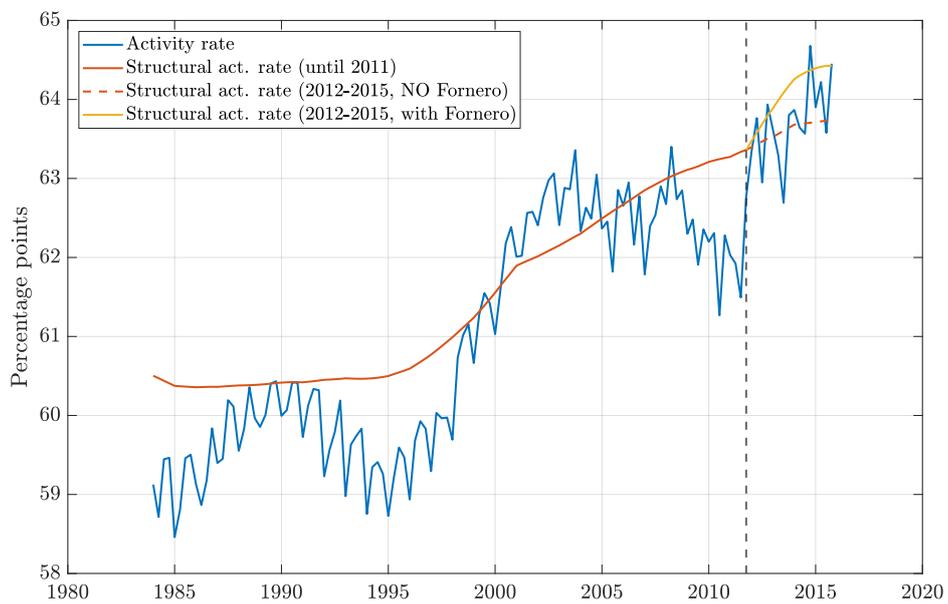


Figure 13: Effect of the Fornero reform on Structural Activity.

Note: The figure plots the quarterly activity rate of the population 15–64 (blue line), the estimated trend activity rate without (red line) and with the reform (yellow line), for the years between 1984q1–2015q4 (see Section 5). The dashed red line represents projections of the trend activity rate until 2015, on the basis of the estimates available in 2011. The red vertical line represents 2011q4, the end of the period prior to the reform.

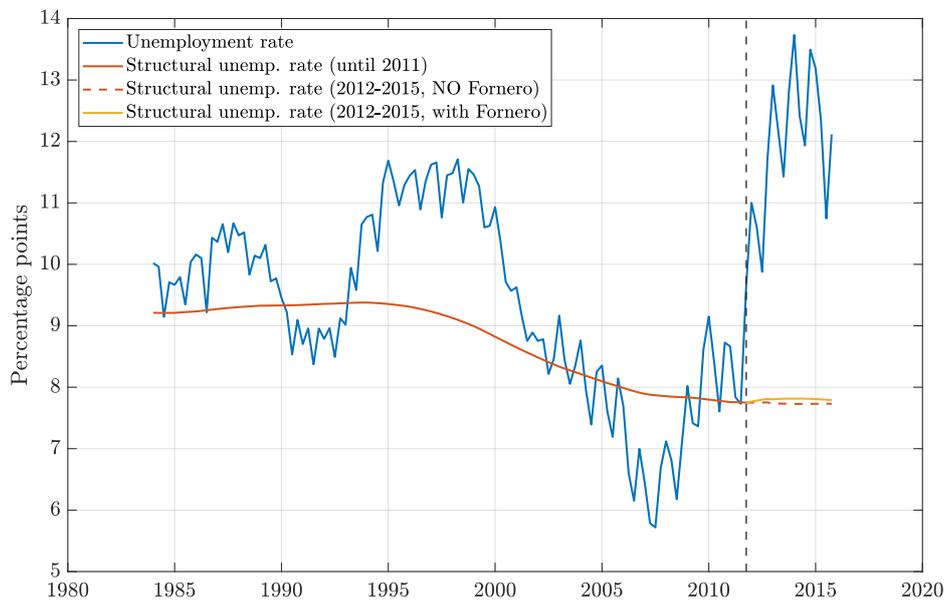


Figure 14: Effect of the Fornero reform on Structural Unemployment.

Note: The figure plots the quarterly unemployment rate of the population 15–64 (blue line), the estimated trend unemployment rate without (red line) and with the reform (yellow line), for the years between 1984q1-2015q4 (see Section 5). The dashed red line represents projections of the trend unemployment rate until 2015, on the basis of the estimates available in 2011. The red vertical line represents 2011q4, the end of the period prior to the reform.



Figure 15: Effect of the Fornero reform on u^*

Note: The figure plots the quarterly unemployment rate (black line) and the projected u^* in the post-reform period (2012q1–2015q4) conditional on trend unemployment rate without the reform (red line) and with the reform (blue-dashed line); see Section 5. Shading denotes the 68% credible interval.

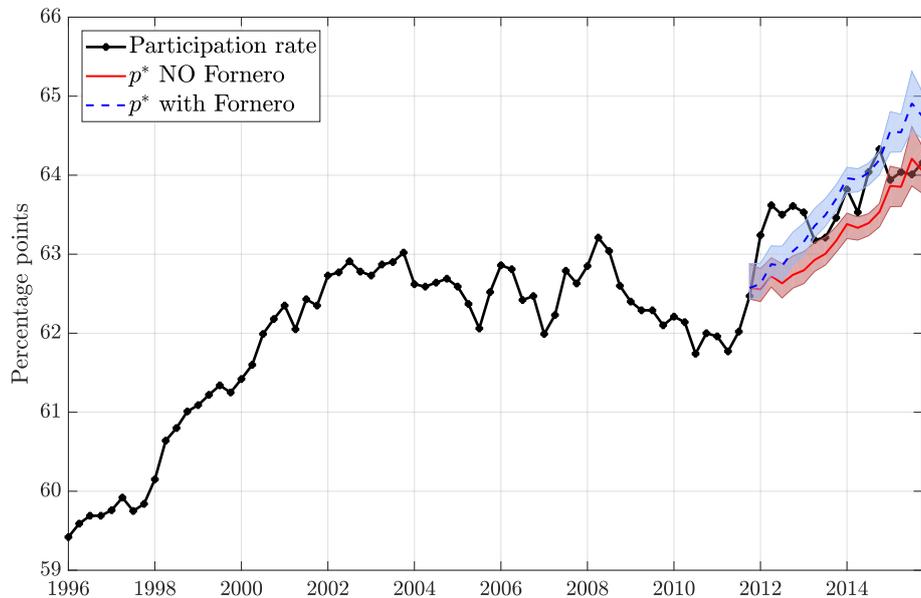


Figure 16: Effect of the Fornero reform on p^*

Note: The figure plots the quarterly participation rate (black line) and the projected p^* in the post-reform period (2012q1–2015q4) conditional on trend unemployment rate without the reform (red line) and with the reform (blue-dashed line); see Section 5. Shading denotes the 68% credible interval.