2

BANCO DE PORTUGAL
ECONOMIC STUDIES

2

Banco de Portugal
Economic Studies

Volume X

Please address correspondence to
Banco de Portugal, Economics and Research Department
Av. Almirante Reis 71, 1150-012 Lisboa, Portugal
T +351 213 130 000 | estudos@bportugal.pt

BANCO DE PORTUGAL
EUROSYSTEM

Lisboa, 2024 • www.bportugal.pt

Banco de Portugal Economic Studies | Volume X – no. 2 | Lisbon, 2024 • **Banco de Portugal** Rua do Comércio, 148 | 1100-150 Lisboa • www.bportugal.pt • **Edition** Banco de Portugal • **Design** Communication and Museum Department | Design Unit • **ISSN** (*online*) 2183-5217

Content

Editor's note[1]

Pedro Duarte Neves

April 2024

---

1. This issue of the Banco de Portugal Economic Studies includes three studies. The first describes the use of robots in Portuguese productive activity. The second highlights the usefulness of surveys on consumer perceptions and expectations of price developments. The third looks at how the human capital accumulated by entrepreneurs affects a possible decision to return to paid work.

2. The opening study of this Banco de Portugal Economic Studies, by Amador, provides a very detailed description of the use of robots in Portuguese productive activity. The main findings of the study are:

(i) The extent to which robots are used in productive activity in Portugal is about half that of the European Union; nevertheless, Portugal is in an intermediate position in this regard (45th percentile);
(ii) The manufacturing industry is, as would be expected, the sector of activity which, in relative terms, has the highest levels of robot usage;
(iii) The use of robots is positively statistically associated, at firm level, with size, productivity, export intensity and profitability;
(iv) The use of robots is negatively associated with the share of labour on income from productive activity, a finding which is in line with empirical literature.

3. The study by Amador shows the possibilities for analysis opened up by the joint use of a wide variety of databases. The first, maintained by the International Federation of Robotics, contains aggregated information on robot usage in productive activity – broken down by sector of activity and type of task performed – for 40 countries and over a period of almost 30 years. The author also uses information provided by Eurostat for EU countries on robot adoption in industry and services. The third source of information on the use of robots, in this case for Portugal, is the "Survey on the Use of Information and Communication Technologies in Enterprises", which makes it possible to identify the extent to which robots are used at firm level. This survey,

---

E-mail: pneves@bportugal.pt

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

along with similar surveys conducted in other European countries, underpins the above-mentioned Eurostat information. Finally, the “Integrated Business Accounts System”, also compiled by Statistics Portugal, provides firm-level information on a number of relevant dimensions, especially in terms of characterising the balance sheet and the presentation of results.

The study by Amador is therefore a very convincing illustration of the gains that can be obtained by simultaneously exploiting databases which, although different in nature and purpose, provide supplementary information on a particular reality.

4. Inflation expectations are central to economic agents’ behaviour and thus to their consumption, savings and investment decisions. They also play a major role in the monetary policy transmission mechanism. As they are not directly observable, central banks use several methods to estimate them: indicators based on surveys of professional market analysts; indicators based on the prices of financial instruments that are indexed to future inflation rates; prospects of economic agents – consumers and businesses (industry, construction, services) – on current and future price developments.

The study by Gomes, Monteiro and Ribeiro looks at euro area consumers’ inflation expectations – over a one-year horizon – and also at perceptions of contemporary price developments. To this end, authors turn to the newly created Consumer Expectations Survey (CES) of the European Central Bank, which covers the period from April 2020 to December 2023 and provides individual information for three characteristics: country, age bracket and income. As a point of reference, the authors also use the European Commission’s Business and Consumer Surveys (BCS), which completed 20 years of existence at the end of 2023.

The main findings of the study are:

(i) Inflation perceptions have a positive correlation with contemporaneous inflation;
(ii) However, consumers’ perceptions of contemporaneous price developments are generally positively biased compared with actual price growth. In this regard, the ECB indicator is more in line with actual inflation than the EC indicator;
(iii) Inflation expectations in one year’s time are statistically linked to perceptions about contemporary price developments and inflation expectations in the preceding month;
(iv) Inflation expectations in one year’s time tend to increase with age and decrease with income level; i.e. older consumers and consumers with lower incomes tend to have higher expectations about future price developments.. This finding, obtained after conditioning on the other variables with predictive value for those expectations, is particularly robust as it was consistently obtained for the six euro area countries considered in this study.

5. The final study of this Banco de Portugal Economic Studies, by Gyetvai, Hossein Dad, Kozeniauskas and Tan, provides empirical evidence on the effects on the income of an entrepreneur who decides to return to paid work after a few years in this role. The results are very clear: such a decision tends to result in large and persistent wage losses, which are larger for longer entrepreneurial experience and higher levels of education. The magnitude of these costs is therefore a factor affecting this type of labour market mobility.

Entrepreneurship is a major driver of economic development through its effects on employment, innovation and competition. This study contributes to a better understanding of this situation in the labour market and may provide a stimulus for developing studies that could better characterise its importance in economic activity in Portugal.

Non-technical summary

April 2024

Robots in Portuguese Firms

João Amador

Robots are one important dimension of the ongoing digital transition. The number of robots in operation around the world has been expanding rapidly. In Portugal, the number of robots per thousand workers increased from 0.6 in 2012 to 1.3 in 2021 (Figure 1). This ratio is much lower than the EU average, which stood at 2.8 in 2021. Germany is the most automated EU economy, with 5.5 robots per thousand workers in 2021.

This study describes the utilization of robots in Portuguese firms. Firstly, it frames Portugal in the international context and describes the distribution of robots along sectors and tasks performed. Secondly, it correlates the existence of robots in the firm with productivity, wages, participation in international trade and profitability. The study uses three databases with different characteristics. The first one, maintained by the *International Federation of Robotics*, contains cross-country aggregate information about

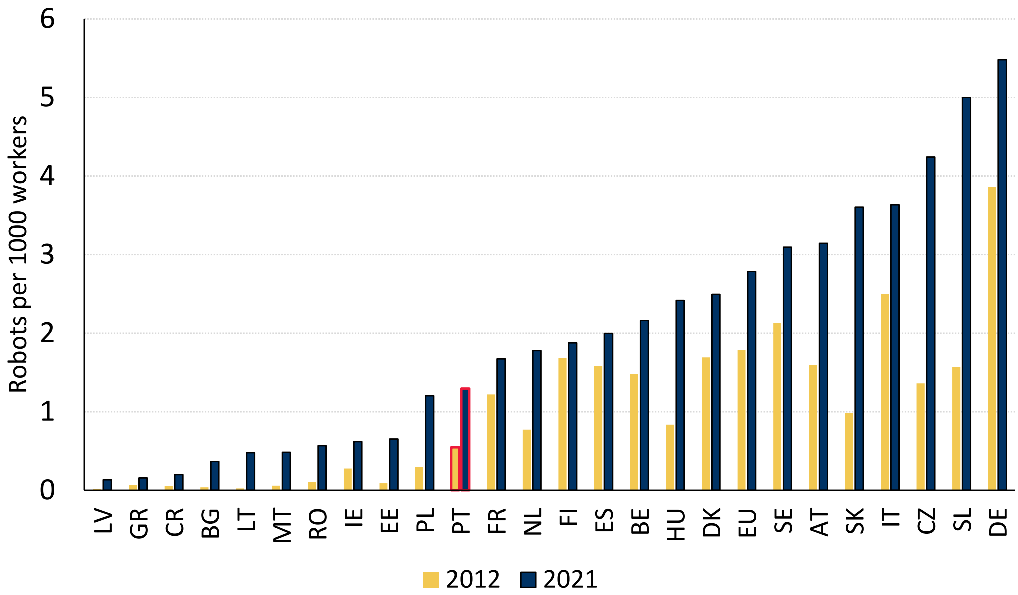

FIGURE 1: Number of robots per worker in the EU

Note: The EU average does not take into account Luxembourg and Cyprus.
Source: World Robotics, AMECO, and author's calculations.

the number, sectors and applications of industrial and services robots in around 40 countries for the period 1993-2021. The other two databases, compiled by Statistics Portugal, contain very rich firm-level information: "*Inquérito à Utilizaçao das Tecnologias de Informação e Comunicação nas Empresas*", which surveys the existence of robots in the firm, and "*Sistema de contas integradas das empresas*", which contains a large number of balance sheet and income statement variables.

Firm-level data confirms that the manufacturing sector is the one with a larger share of robots, not only in number of firms, but also when weighting with turnover and employment. As for the association between the existence of robots at the firm and its performance, results show a positive relationship with productivity, export intensity and profitability, while the correlation with the labour share is negative. The impact on wages is not statistically different from zero. Moreover, the positive correlation between robots and productivity is higher in the lower quantiles of the productivity distribution, i.e., the productivity gain associated to the existence of robots is larger in the group of less productive firms.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. X, No. 2

Robots in Portuguese Firms

João Amador
Banco de Portugal
Nova SBE

April 2024

Abstract
Robots are one important dimension of the ongoing digital transition, even if they are not a truly recent technology. The discussions on the fastly growing number and sophistication of robots, as well as their impacts on jobs and wages are abundant but arguments often lack empirical evidence. This article tries to contribute to this research strand by presenting evidence on the prevalence of robots in Portuguese firms and correlating their presence with performance indicators, such as productivity, wages, international trade and profitability. We conclude that Portugal is lagging the EU in terms of robots' adoption, which are associated with higher productivity, lower labour share, higher export intensity and higher profitability. (JEL: O14, O3, D24, D33)

Keywords: Robots, productivity, wages, profits, trade, Portugal, firm-level data.

1. Introduction

Robots are one dimension of the technological transformation that has been taking place in the last decades. Robots are very diverse and have captured the collective imagination mostly due to the fear of a large scale replacement of humans in the labour market, leading to job destruction and unemployment. Such fears are recurrent whenever new labour-replacing technologies arrive to the production process and jobs are destroyed. The terms "automation" and "robotics" are often used interchangeably, but have different meanings. Automation is the process of using technology to perform human tasks while robotics concerns developing machines to carry out a particular function. In this perspective, not all types of automation use robots and not all robots are designed for process automation. That said, most robots are used for process automation, especially in manufacturing. The manufacturing sector has long been at the forefront of robots' usage to assist production. Automated mobile robots (AMRs), automated guided vehicles (AGVs) and articulated robots are deployed on

Acknowledgements: The author is thankful to Nuno Alves, António Antunes, Joana Garcia, Nicholas Kozeniauskas, Pedro Duarte Neves and Ana Catarina Pimenta for their very useful comments and suggestions. The opinions expressed are those of the author and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.

E-mail: jamador@bportugal.pt

factory floors and in warehouses to fasten processes, increase efficiency and promote safety. They are used across a variety of applications, including handling, welding and soldering, assembling and dispensing. These technologies are also increasingly used in agriculture, transportation, logistics, and healthcare. Therefore, robots are very diverse and, contrary to common wisdom, they are not necessarily extremely sophisticated machines.

The literature on the economic impact of robots on firms is too vast to be listed here, but the impacts on productivity, wages, labour share, employment, inequality and exports, all of them depending on firms' characteristics and skills, dominate the list of research questions. Most studies are empirical and use firm-level data, which in some cases is rich enough to prove causal relationships.

Autor (2015) has set the stage for many of these discussions, stating that the past two centuries of automation and technological progress have not made human labor obsolete. It argues that automation complements labor, raises output in ways that lead to higher demand for labor, and interacts with adjustments in labor supply. The paper also draws attention to the fact that wage gains went disproportionately to those at the top and at the bottom of the income and skill distribution, though also arguing that this in unlikely to continue very far into the future. The conceptual debate continued with Acemoglu and Restrepo (2018), which presents a framework to study the implications of automation and AI on the demand for labor, wages, and employment. The framework signals a displacement effect coming from the replacement of labor in several tasks. This effect reduces the demand for labor and wages, but it is counteracted by a productivity effect, which increases the demand for labor in non-automated tasks. Such productivity effect is complemented by additional capital accumulation and more automation, further increasing the demand for labor. Overall, automation increases output per worker more than wages and reduces the share of labor in national income. The paper also highlights possible constraints and imperfections that may slow down the adjustment of the economy and the labor market, for example a mismatch between the skill requirements of new technologies and existing ones.

The debate evolved towards optimal policies to implement in this context. One contribution is Guerreiro *et al.* (2021), which uses a quantitative model that features technical progress in automation and endogenous skill choice to show that, given the current U.S. tax system, a sustained fall in automation costs can lead to a massive rise in income inequality. Authors argue that it is optimal to tax robots while the current generations of routine workers, who can no longer move to non-routine occupations, are active in the labour force. Once these workers retire, optimal robot taxes are zero.

The cross country dimension of the empirical analysis is covered by de Vries *et al.* (2020), which examines the impact of industrial robots on jobs by combining data on robot adoption and occupations by industry in 37 countries for the period 2005-2015. The article concludes that a rise in robot adoption relates significantly to a fall in the employment share of routine manual task-intensive jobs. This relation is observed in high-income countries, but not in emerging market and transition economies. At the sub-national level Acemoglu and Restrepo (2020) study the effects of industrial robots on US labor markets and show theoretically that they may reduce employment and wages.

The paper uses variation in industry-level advances in robotics and local industry employment to confirm the negative effects of robots on employment and wages across commuting zones. The paper also reports that robots' impact is distinct from other types of capital and technologies.

Acemoglu *et al.* (2020) studies the implications of robot adoption in France. Authors conclude that adopters experienced significant declines in labor share, the share of production workers in employment, and increases in value added and productivity. The observed expansion in firms' employment comes at the expense of competitors, leading to an overall negative association between adoption and employment. Moreover, robot adoption has a large impact on the labor share because adopters are larger and grow faster than their competitors. Aghion *et al.* (2020) also uses micro data for the French manufacturing sector between 1995 and 2017 to assess the causal effects of automation technologies on employment, sales, prices, wages, and the labor share. The estimated impact of automation on employment is positive, even for unskilled industrial workers, it leads to higher sales, higher profits, and lower consumer prices, while it leaves wages, the labor share and within-firm wage inequality unchanged. The authors refer that, in a globalized world, attempts to curb domestic automation in order to protect domestic employment may be self-defeating due to foreign competition. Still with French data, Bonfiglioli *et al.* (2020) studies how imports of industrial robots affect firm-level outcomes. Robot importers are larger, more productive, and employ a higher share of managers and engineers. Authors conclude that robot imports increase productivity and the employment share of high-skill professions, but have a weak effect on total sales.

More recently, Deng *et al.* (2023) analyses the impact of robot adoption on employment composition using data on robot use in German manufacturing plants linked with social security records and data on job tasks. The event-study analysis of robot adoption signals more favourable employment effects for the least routine-task intensive occupations and for young workers, with the latter being better at adapting to change. Authors conclude that the displacement effect of robots is occupation biased but age neutral, whereas the reinstatement effect is age biased and benefits young workers most.

As for the international trade dimension, Alguacil *et al.* (2022) studies the causal effect of robot adoption on the intensive and extensive margins of exports of Spanish manufacturing firms over the period 1990–2014, and finds that those adopting robots experience a sharp increase in export probability, exports and share of exports in total output. The positive impact of robot adoption on exports is driven by its positive effect on firm total factor productivity (TFP), product innovation and imports. Also for Spanish manufacturing firms, Koch *et al.* (2021) uses a rich 27-year long panel data set and provides causal evidence on the characteristics that prompt robot utilization and their impact on adopting firms relative to non-adopting firms. The article concludes that there is robust evidence for positive selection, i.e., firms more productive ex ante are more likely to adopt robots and, conditional on size, firms more skill-intensive ex ante are less likely to do so. In addition, robot adoption generates substantial output gains, reduces the labour cost share, and leads to net job creation. The results also demonstrate a positive causal effect of robots on productivity, as well as a complementarity between robots and foreign sales in increasing productivity. In addition, Leone (2022) uses a panel of Spanish manufacturing firms covering the period 1990-2017 and shows that those acquired by multinationals start investing in industrial robots, which leads to a reduction of the labor share at the firm and industry levels. The estimates of the model imply that, without multinationals and robots, the manufacturing labor share would be at the level of two decades ago.

As for Portugal, studies about robot adoption are Candeias *et al.* (2022), which studies the implications of automation on productivity and employment in the automotive sector, and Almeida and Sequeira (2023), which uses a fixed effects panel quantile regression and an instrumental variable regression model to study the impact of robots, software, ICT and physical capital on productivity. Amador and Silva (2023) studies the association between new digital technologies, including robots, and productivity, wages and export intensity.

In this article we describe the prevalence of robots in Portuguese firms. Firstly, we frame Portugal in the international context and describe the distribution of robots along sectors and tasks performed. Secondly, we test regressions with controls for sector and firm size to correlate the existence of robots at the firm with productivity, wages, international trade and profitability. The short time-span of the data does not allow to study employment. The article concludes that Portugal has been increasing the number of robots in operation but it starts from very low levels and other European laggards have evolved much faster. Moreover, firms with robots are also those with higher productivity, profitability and export intensity, while the labour share is lower. Moreover, the positive association between robots and productivity is higher in the lower quantiles of the productivity distribution.

The article is organized as follows. In section Section 2 we present the three databases used in the article. Section 3 presents cross country data on the number and type of robots since the early nineties to frame the Portuguese situation in the international context, while also presenting evidence for the US and China. Section 4 explores the sectoral dimension of robots' adoption by Portuguese firms. Section 5 is divided in two parts, the first concerning the distribution of selected performance variables for firms with robots versus those without, and the second presenting a correlation exercise based on a regression with time and sector fixed effects. This exercise is deepened by including a quantile regression. Finally, Section 6 presents some concluding remarks.

2. Data

In this article we use three main databases with different characteristics. The first one contains cross-country aggregate information about the number, sectors and different applications of industrial and services robots in around 40 countries for the period 1993-2021 (International Federation of Robotics (2022)). The database is maintained by the *International Federation of Robotics*, a non-profit organization established in 1987, whose members come from the robotics industry, national or international industry associations and research & development institutes from more than 20 countries.

The other two databases contain rich firm-level information and we merge them to assess heterogeneity in robot utilization by Portuguese firms and to establish an association between the existence of robots at the firms and productivity, labour market outcomes, profitability and participation in international trade. The first set of data contains firm's answers to "*Inquérito à Utilizaçao das Tecnologias de Informação e Comunicação nas Empresas*" (IUTICE), a survey conducted by the Portuguese national institute of statistics (Statistics Portugal). This statistical operation is carried out annually within the framework of EU legislation (EC regulation No. 808/2004), which establishes a set of harmonization guidelines, thus ensuring the availability of comparable statistical results across member states. This is the set of data underlying the computation of aggregate digitalization statistics by the Eurostat. The IUTICE was initiated in 2003 and we use information until 2020. The set of firms surveyed is not constant and the size of the sample has changed along the years, with a notable increase after 2010, which improved its representativeness. The set of questions posed to firms has changed substantially along the different vintages of the survey. The editions of 2018 and 2020 include an explicit question regarding the existence of robots at the firms, which are the core of our information. The short time span of the data turns some analysis non viable, notably the assessment of changes in the number of employees at the firm after the adoption of robots, not to mention proofing causal inference.

Another very relevant issue, with potential strong impact on the results, is the non-quantification of the number of robots in operation in the firm. Indeed, to have just one robot performing a particular task is quite different from having an entirely automated production line. A proxy for the actual number of robots operating in firms could be obtained from very detailed firm-product data on imports, more specifically from the accumulation of imports of goods classified as robots. Nevertheless, this firm-level data source is not possible to merge with other firms' characteristics and robots are not necessarily imported by the firms that incorporate them in the productive process.

The second firm-level database used is the "*Sistema de contas integradas das empresas*", compiled by Statistics Portugal. This database builds on mandatory legal reporting by Portuguese firms to Statistics Portugal, tax administration, Banco de Portugal and Ministry of Justice. It covers virtually the universe of Portuguese firms, including self proprietorships. This dataset contains a large number of balance sheet and income statement variables, which allow us to control for firm heterogeneity and to compute labour productivity (GVA per worker) and TFP. Merging the two firm-level datasets is straightforward since there is a common firm identifier.

3. Time and country comparisons

The number of robots in operation around the world has been expanding rapidly. In this section we look at the Portuguese experience with robots from an aggregate and temporal perspective, and compare with the other EU countries, US and China. In addition, comparing not only the number of robots but also their type is informative.

In order to compare internationally, the number of robots in the country must be standardized relatively to some measure of size. Panel A) of Figure 1 plots the number of robots per thousand million of euros of GDP, at constant prices of 2015. From 1995 to 2021, the ratio increased almost ten-fold in Portugal, and it is at around three quarters of the EU average. Although also increasing in this period, the ratio for the US is one third of the EU average. Strikingly, China displays a very fast increase in this period and the level of robots on GDP is twice larger than in the EU in 2021. Panel B) of Figure 1 sets the total employment as the standardization variable and results are somewhat different due to distinct labour productivity levels in these countries. In Portugal, the number of robots per thousand workers increased from virtually zero in 1995 to 1.3 in 2021. In this final year, the ratios for the US, the EU average and China were 2.1, 2.8 and 1.7, respectively. Germany clearly stands out as a highly automated economy, with 5.5 robots per thousand workers in 2021 and a steady increase since 1995.

The very high number of robots per worker in Germany versus other economies is also visible in Figure 2, which ranks the ratio for 25 EU countries. Portugal is placed in the mid-lower end of the distribution in 2021 in a context where several countries that were lagging in 1995 recorded very strong progress (e.g. Czech Republic, Slovenia, Slovakia).

Table 1 presents the distribution of robots along sectors in 2012 and 2021 in Portugal, Germany, EU, US and China. In the first of these two years the share of "unspecified" is large. Despite the limitations of the database, it is clear that manufacturing is dominant in terms of robots' usage in all these countries, with shares above 80 per cent (close to 90 per cent in Portugal). Regarding robot utilization, the automotive sector is the largest manufacturing sector in Portugal, Germany, EU and US, but in China it weights less than the electrical and electronics sector. In Portugal, the sector of rubber and plastics and the non-auto metal products rank second and third, with shares of 12.4 and 11.1 per cent in 2021, respectively.

	Portugal		EU		Germany		US		China	
SECTOR	2012	2021	2012	2021	2012	2021	2012	2021	2012	2021
Manufacturing	75.8	89.8	87.7	81.6	89.6	82.4	70.9	86.5	67.3	80.5
Food and beverages	3.7	4.3	5.2	6.2	3.7	2.8	3.4	5.7	1.2	1.7
Rubber and plastic (non-auto)	4.7	12.4	8.9	8.0	7.7	7.3	4.9	4.9	11.5	3.5
Metal products (non-auto)	22.1	11.1	8.2	8.6	5.5	6.7	5.4	2.7	2.7	4.5
Industrial machinery	1.9	5.2	4.6	6.3	4.3	5.5	1.0	1.6	0.7	5.9
Automotive	36.8	47.1	45.6	39.6	53.3	48.8	39.2	41.6	33.3	26.3
Electrical and electronics	2.1	1.0	3.9	4.0	4.9	4.8	12.3	15.3	10.8	30.1
Other sectors	1.9	0.9	1.3	2.3	1.4	1.4	0.3	1.2	0.3	2.0
Unspecified	22.3	9.2	11.0	16.1	9.0	16.1	28.8	12.3	32.3	17.4
Total	100	100	100	100	100	100	100	100	100	100

TABLE 1. Distribution of robots in the economy, by sector of activity (2012 and 2021)

Note: The EU average does not take into account Luxembourg and Cyprus.
Source: World Robotics and author's calculations.

Table 2 complements the previous analysis with the breakdown of robots in Portugal, Germany, the EU average, US and China in 2012 and 2021, along the different types of tasks they perform. As in the other countries considered, robots performing handling

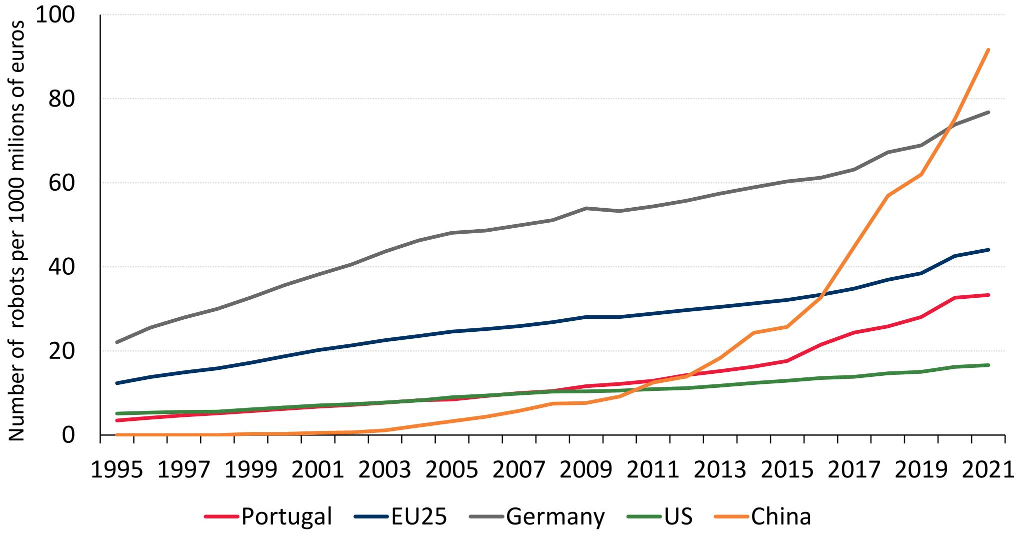

(A) Robots on GDP (constant prices of 2015)

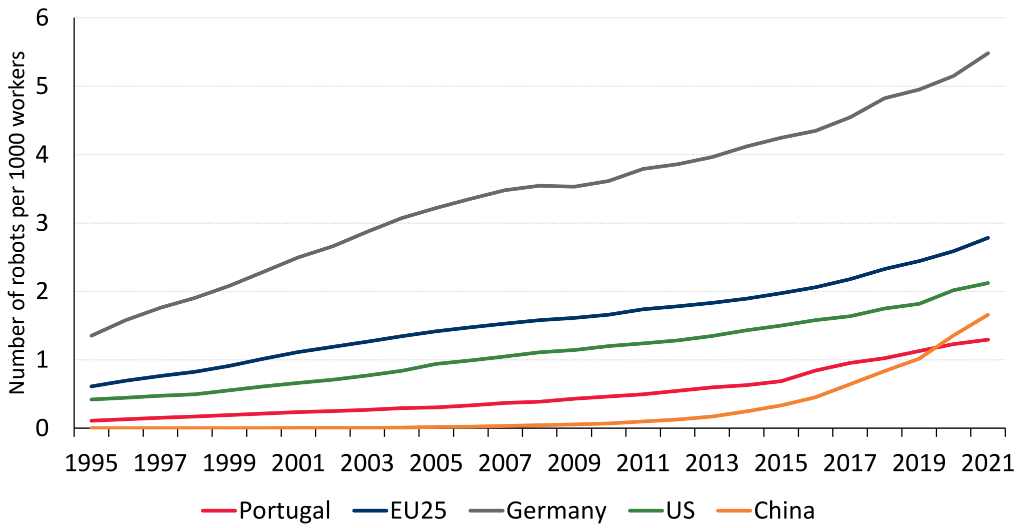

(B) Robots per worker

FIGURE 1: Robots across the world

Note: The EU average does not take into account Luxembourg and Cyprus.
Source: World Robotics, AMECO, Penn World Tables and author's calculations.

and machine tending operations are dominant in Portugal in 2021 (51.9 per cent), though in 2012 welding and soldering played the largest role (46.2 per cent). The remaining tasks are clearly less important in all countries, with the notable exception of robots for assembling and disassembling in China, with a share of 15.9 per cent in 2021.

Information regarding the utilization of robots in firms must be collected through surveys or administrative data. The Eurostat collects data from standardized national surveys on adoption of digital technologies by households and firms and provides

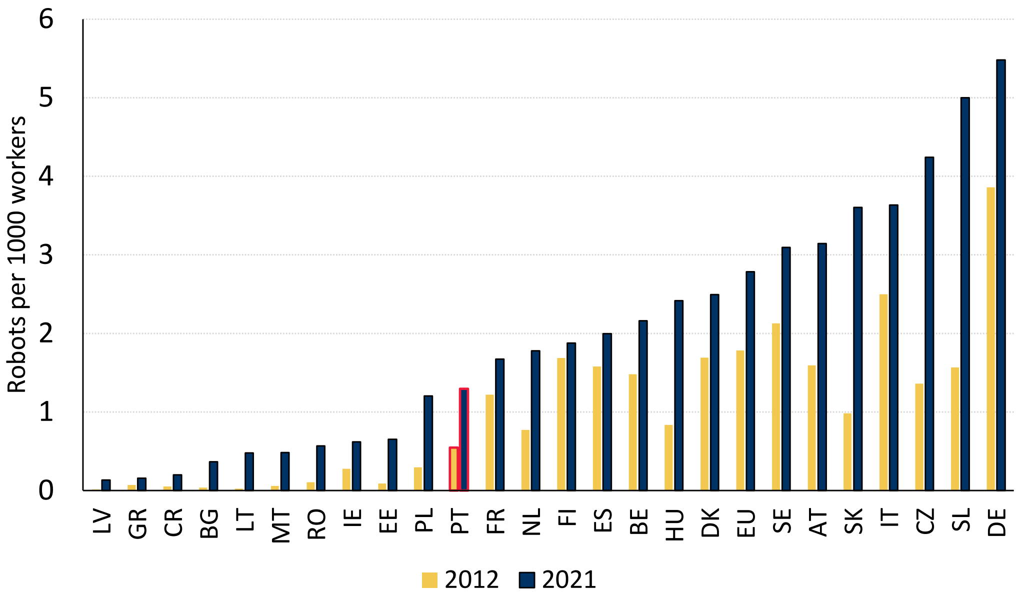

FIGURE 2: Number of robots per worker in the EU

Note: The EU average does not take into account Luxembourg and Cyprus.
Source: World Robotics, AMECO and author's calculations.

	Portugal		EU		Germany		US		China	
TASK PERFORMED	2012	2021	2012	2021	2012	2021	2012	2021	2012	2021
Handling and machine tending	27.7	51.9	52.6	58.1	53.8	56.4	36.2	44.9	32.9	40.4
Welding and soldering	46.2	25.3	25.8	17.0	24.4	15.9	37.3	26.2	43.7	25.6
Dispensing	2.2	2.8	4.2	3.2	4.8	4.0	4.5	3.2	6.0	3.7
Processing	1.6	1.4	3.3	2.7	2.8	2.8	2.2	1.1	1.2	1.6
Assembling and disassembling	2.7	4.4	5.6	5.7	5.9	5.8	7.8	6.3	7.3	15.9
Others	0.5	6.6	2.1	2.7	1.7	2.6	6.6	6.3	2.9	3.2
Unspecified	19.1	7.5	6.4	10.7	6.6	12.5	5.3	11.9	6.0	9.5
Total	100	100	100	100	100	100	100	100	100	100

TABLE 2. Distribution of robots in the economy, by tasks performed (2012-2021)

Note: The EU average does not take into account Luxembourg and Cyprus.
Source: World Robotics and author's calculations.

online a quite large set of aggregate results.[1] Even if the survey only captures the existence of robots and not how many of them actually operate in the firm, it is still very useful information. Panel A) of Figure 3 ranks the percentage of firms with robots in the EU member countries in 2020. Denmark ranks first with almost 20 per cent of firms reporting the existence of robots. Portugal lays in the mid-low segment of the distribution, in the 18th position, with a share slightly above 8 per cent. Panel B) of

1. This database by the Eurostat is available at https://ec.europa.eu/eurostat/web/digital-economy-and-society/database/comprehensive-database.

Figure 3 presents the share of firms with services versus industrial robots in 2020.[2] With the exception of Ireland, the share of firms using industrial robots is higher than that of those using service robots. The share of those using service robots in Portugal is about half of those using industrial robots, even if some firms may have both types. The sum of the shares of firms with industrial and services robots in a country in panel (B) may exceed the total in panel (A) because some firms have both types of robots in operation.

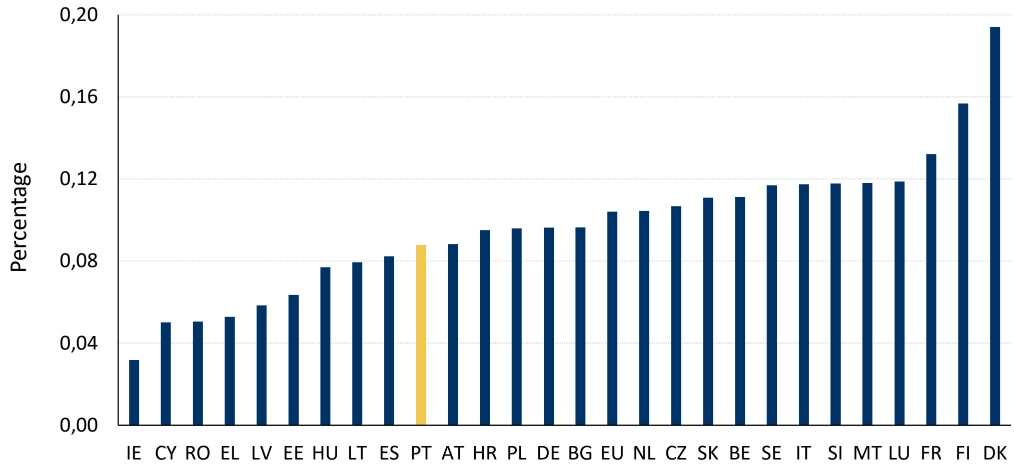

(A) Total robots

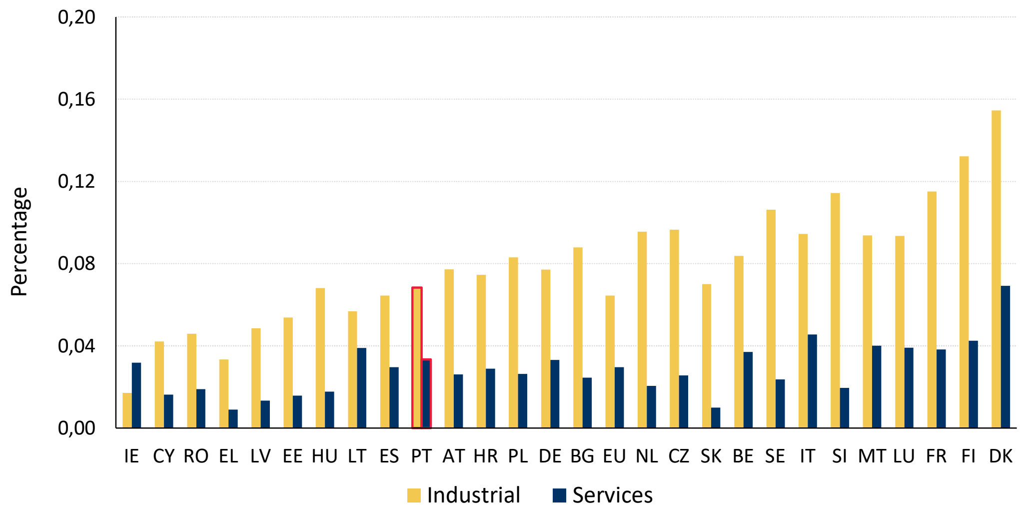

(B) Industrial and services robots

FIGURE 3: Percentage of firms with robots in EU countries in 2020
Source: Eurostat and author's calculations.

2. Service robots typically operate in sectors such as healthcare, hospitality, retail, and logistics. They execute functions ranging from cleaning and delivery to care-giving and customer service.

4. Robots within sectors in Portugal

In this section we use the firm-level data to analyse the share of firms that report the existence of robots within the 12 main sectors in the economy, weighting firms surveyed according to population, sales and employment. Complementarily, we take a NACE 2-digit classification of activity and rank sectors according to the percentage of firms with robots within each sector using population weights.

Table 3 shows again that manufacturing is the sector where robots are most pervasive. This is true in terms of number of firms and also when they are weighted according to sales and employment. When we use these weights, robots are also quite pervasive in transports and other services, and at a distance, in the sectors of construction and water.

SECTOR	Number	Turnover	Employment
Manufacturing	11.1	48.3	33.9
Electricity & gas	4.5	1.1	10.4
Water	5.9	14.1	16.5
Construction	3.1	17.0	10.8
Wholesale & retail	1.6	5.2	4.0
Transport	1.7	23.3	20.6
Accommodation	1.0	1.0	1.1
Information & com.	1.0	0.5	0.8
Real estate	0.6	0.4	0.2
Consult. & science	0.7	1.2	1.0
Administrative act.	1.2	1.2	2.2
Other services	2.3	24.8	19.3

TABLE 3. Share of firms with robots along sectors (2018)

Note: Sectors of Agriculture, Education, Health & social and Arts & sports are not reported due to the small number of responding firms.

The number of firms surveyed and included in our database (5964 in 2018 and 5382 in 2020) is large enough to cover the largest sectors of activity and ensure a reasonable representativeness, but that is not necessarily the case if we take a very detailed sectoral classification. Even so, aiming at a finer picture of robot utilization, we check the percentage of firms reporting the existence of robots in operation in 2018, weighting each one of them by the respective turnover in the sector. Figure 4 ranks sectors and labels the top ten. Among these top ten sectors, two that are typically associated with strong automation emerge: manufacture of motor vehicles and production of metal products, except machinery and equipment. Activities associated with transport are also represented (air transport and postal and courier), as well as those of construction (civil engineering and construction of buildings). In addition, 25 per cent of firms operating in the sector of repair of computers and household goods report the existence of robots in operation.

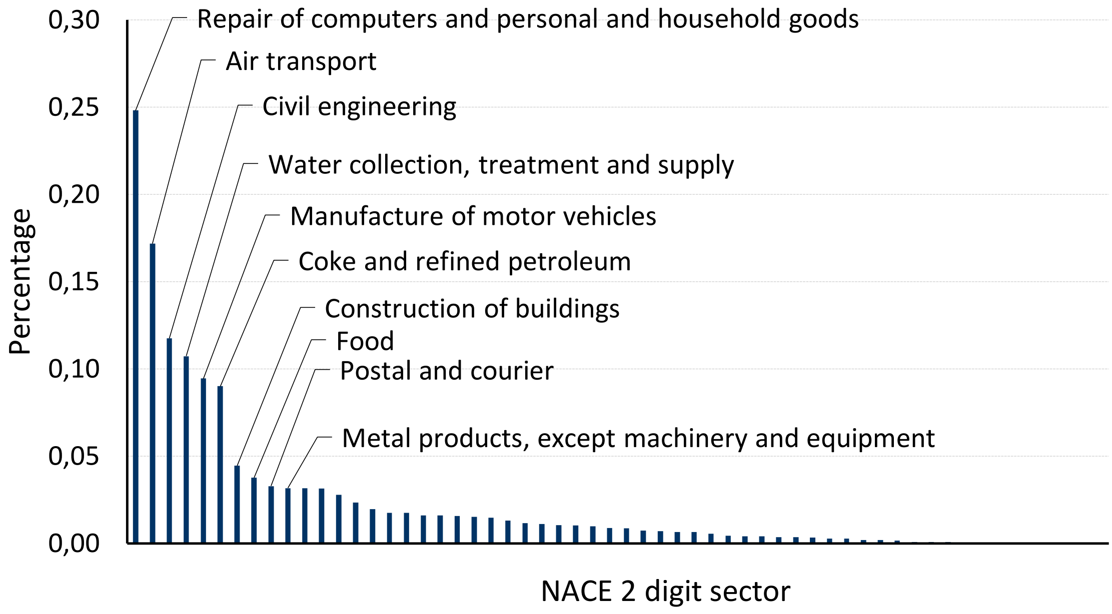

FIGURE 4: Robots per sector
Note: Individual firms are weighted according to their turnover.

5. Robots, productivity, wages, trade and profitability

In this section we move one step further and study the correlation between the existence of a robot at the firm and its performance in several dimensions: productivity, wages, international trade and profitability. We start by comparing the distribution of performance variables for firms with and without robots. Next, we run a set regressions with size, and year and sector fixed-effects to assess these relationships in a statistically robust way.

5.1. Distributions

Panel A) of Figure 5 plots the kernel distributions of TFP for the group of firms that report the existence of a robot versus those that report the absence of this type of machinery. We compute the firm-level TFP according to the method developed by Levinsohn and Petrin (2003). The procedure was implemented using the STATA command "prodest", which estimates the production functions using a control function approach. By default, the command requires the log gross output variable (in our case, the log of the GVA, at market prices), a set of free variables (typically the log of labor), a set of state variables (the log capital) and a set of proxy variables (in our case, the cost of goods sold). The capital stock corresponds to total fixed assets of the firm, as reported in the balance sheet. It should be noted that the inclusion of robots as an input to the production function, which would allow a more direct reading of their effects, was not considered because the database only signals the existence of robots and not their number.

The TFP distribution of firms with robots is clearly shifted to the right when compared with that of non-adopters. Its skewness is negative, as opposed to a positive one for firms without robots. When we compare the kernel distributions for the logarithm of labour productivity, in Panel B) of Figure 5, the conclusion that firms with robots are also those with higher productivity emerges again. These initial results are compatible with the studies that indicate positive impacts of robots on firms' productivity.

Panel A) of Figure 6 plots the kernel distributions for the logarithm of wages in the group of firms that refer the existence of a robot versus those that do not report having them. As referred in empirical studies for other countries, and consistent with their higher productivity, wages are higher in firms with robots. Panel B) of Figure 6 repeats the comparison but focusing on the labour share. The distribution of firms reporting the existence of robots is shifted to the left, signalling a lower labour share, which is also in accordance with results obtained for other countries.

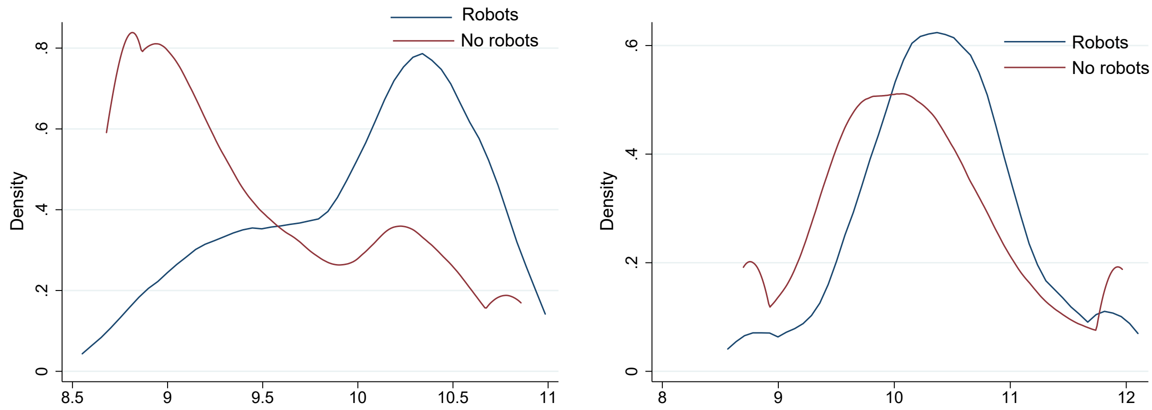

(A) TFP (B) Labour productivity (ln)

FIGURE 5: Productivity indicators

Next, we turn to international trade. Panel A) of Figure 7 plots the kernel distributions for the export intensity ratio, measured as exports on turnover, in the group of firms that refer the existence of a robot versus those that do not report having them. Both distributions present the usual bimodal shape but firms that report having robots post much higher densities for larger export intensities, thus signalling stronger export activity. As for imports, which are equally important as regards trade integration and gains, panel B) of Figure 7 shows that firms with robots are also those with larger shares of imports on turnover.

Lastly, we look at differences in the distributions of profitability indicators for adopters and non-adopters of robots. Panel A) of Figure 8 shows the two distributions for the ratio of earnings before interest, taxes, depreciation and amortization (EBIDTA) on total assets, while panel B) focuses on the ratio of gross operating surplus on GVA.

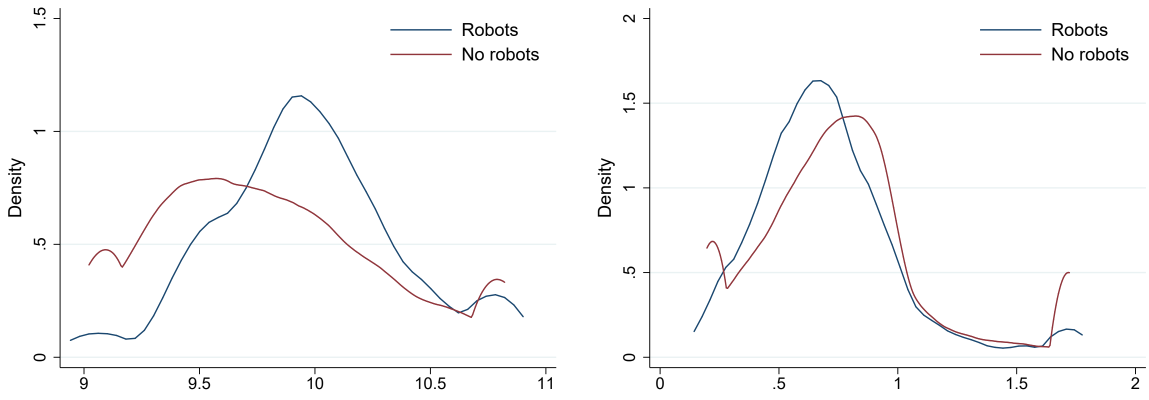

(A) Wages (ln) (B) Labour share

FIGURE 6: Wage indicators

Note: In order to correct for extreme observations, variables were subject to a winsorization procedure affecting obsercations up to the percentile 5 and above percentile 95.

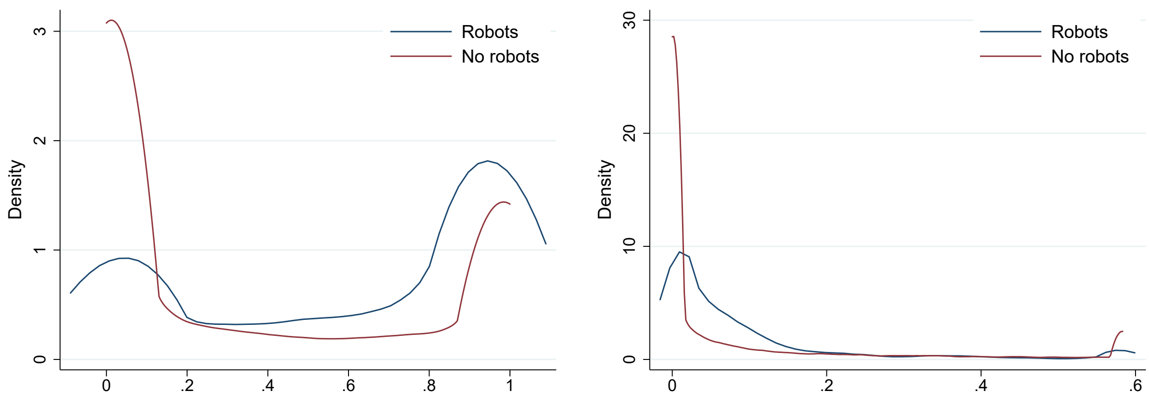

(A) Export intensity (B) Import intensity

FIGURE 7: International trade indicators

Note: In order to correct for extreme observations, variables were subject to a winsorization procedure affecting obsercations up to the percentile 5 and above percentile 95.

In both panels, the differences between distributions are small but the distributions for firms that report the existence of robots have more density in higher profitability ratios.

5.2. *Correlations*

The previous subsection signals correlations by visually comparing distributions, but this is quite partial and lacks statistical support. Although having two years of data does not allow us to test causality or to consider time-invariant characteristics of the

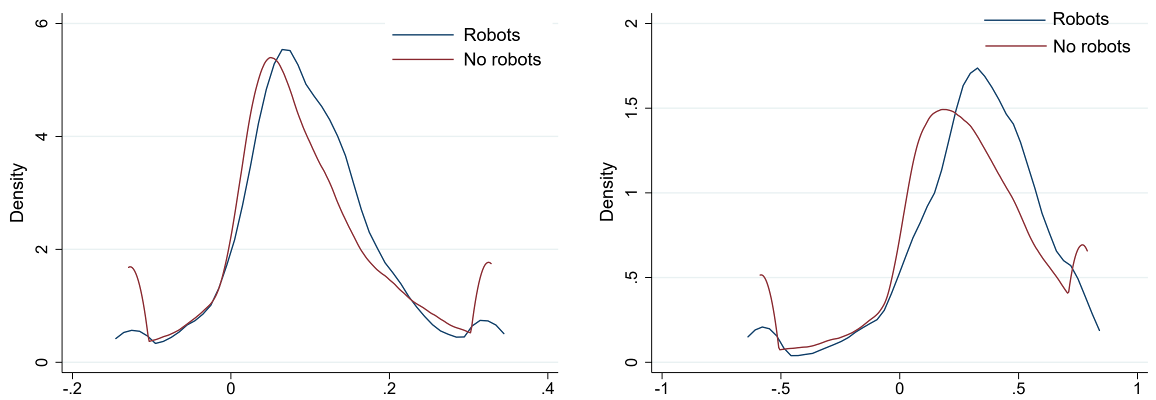

(A) EBITDA on assets (B) Gross operating surplus on GVA

FIGURE 8: Profitability indicators

Note: In order to correct for extreme observations, variables were subject to a winsorization procedure affecting obsercations up to the percentile 5 and above percentile 95.

firms that might affect the dependent variable, it is still possible to run regressions that take into account time and sector fixed effects and control for firm size.

Table 4 presents the results of the regression exercise, considering the eight dependent performance variables separately, with a dummy variable that takes the value one if the firm reports the existence of a robot and a control for the firm size class. The four size classes correspond to micro, small, medium and large firms, in accordance with the definition used by the European Commission, which combines turnover and number of employees. Regressions include time and sector fixed effects, according to the 2-digit NACE classification of economic activities.

Columns one and two of table 4 post positive and significant coefficients in the dummy variable, signalling a positive correlation between the existence of robots at the firm and both TFP and labour productivity, which is in accordance with what is found in the literature for other countries. As for the correlation between robots and the labour share (column 3), the coefficient is negative and significant, while the correlation with the logarithm of wages is not statistically different from zero (column four). In what concerns international trade, the correlation of robots with the export intensity ratio is positive but not statistically different from zero in the import intensity. Finally, as for EBIDTA on assets and gross operating surplus on GVA, the coefficient of the dummy variable is positive, though slightly less significant in the former profit ratio. The coefficients of the size class variable in the regressions are also in accordance with the literature. Larger firms are more productive, accordingly they pay higher wages, they post a lower labour share, they are more engaged in exports and they are more profitable.

Results in Table 4 can be complemented by estimating the coefficients on the impact of having robots at the firm in different segments of the distribution of the outcome variables. We adopt a quantile regression approach that allows for the inclusion of

VARIABLES	(1) TFP	(2) Labour productivity	(3) Labour share	(4) Wages	(5) Export intensity	(6) Import intensity	(7) EBITDA over assets	(8) GOS on income
Robots	0.077***	0.044**	-0.043***	-0.008	0.098***	0.487	0.007**	0.040***
	(0.010)	(0.021)	(0.010)	(0.014)	(0.011)	(0.603)	(0.003)	(0.010)
Size class	0.428***	0.207***	-0.025***	0.193***	0.125***	-0.238	0.007***	0.023***
	(0.003)	(0.005)	(0.002)	(0.004)	(0.003)	(0.423)	(0.001)	(0.002)
Constant	8.610***	9.752***	0.819***	9.405***	0.084***	1.138	0.077***	0.199***
	(0.006)	(0.011)	(0.006)	(0.009)	(0.007)	(1.060)	(0.002)	(0.005)
Nb obs.	17,228	17,228	17,228	17,114	13,643	13,643	17,228	17,228
Adj. R^2	0.735	0.360	0.153	0.346	0.447	0.007	0.055	0.125
Time FE	YES	YES	YES	YES	YES	YES	YES	YES
Sector FE	YES	YES	YES	YES	YES	YES	YES	YES

TABLE 4. Existence of robots at the firm and performance

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Robots is a dummy variable assuming the value 1 in case the firsms reposts the existence of robots. EBITDA stands for earninsg before interest, taxes, depreciations and amortizations. GOS stands for gross operating surplus. In order to correct for extreme observations, all dependent variables were subject to a winsorization procedure affecting observations up to the percentile 5 and above percentile 95, except in the cases of TFP and labour productivity, where the percentiles are 1 and 99.

fixed effects, as proposed by Machado and Santos Silva (2019) and implemented in the STATA package MMQREG (Rios-Avila (2020)). The database and the sets of explanatory variables and fixed effects are exactly the same as in Table 4.

The eight panels of Figure 9 report the coefficients for the dummy variable that signals the existence of robots at the firm along the distribution of each outcome variable. Panels (A) and (B) show that the positive correlation between robots and productivity is larger in the lower quantiles, i.e. for less productive firms. In can be argued that firms on the top of the productivity distribution have already explored several margins of efficiency, thus the marginal benefit of robot adoption is lower. In addition, the positive correlation of robots and the export intensity of the firm is weaker for those with high exports on sales (panel E). As for the remaining outcome variables, given the amplitude of the confidence bands, there is no statistically significant profile for the coefficients in the different percentiles.

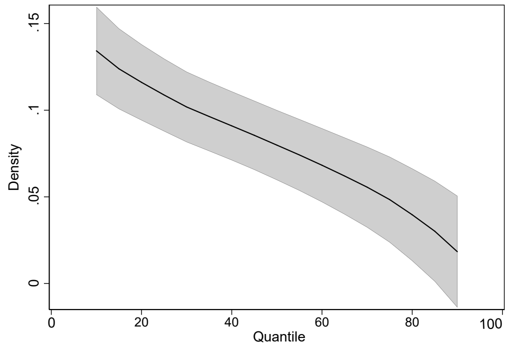

(A) TFP

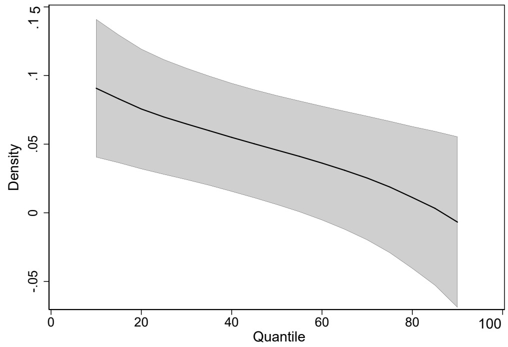

(B) Labour productivity (ln)

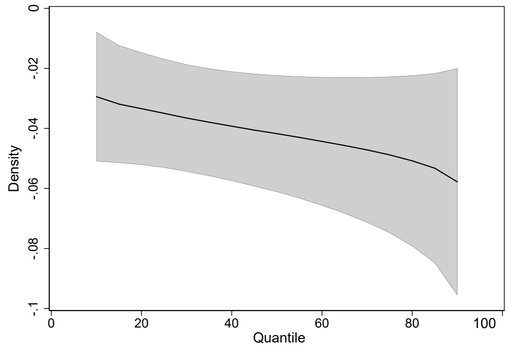

(C) Labour share

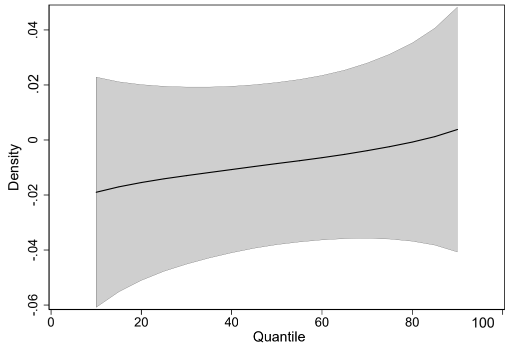

(D) Wages (ln)

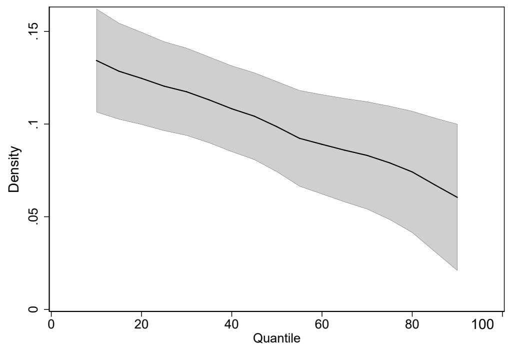

(E) Export intensity

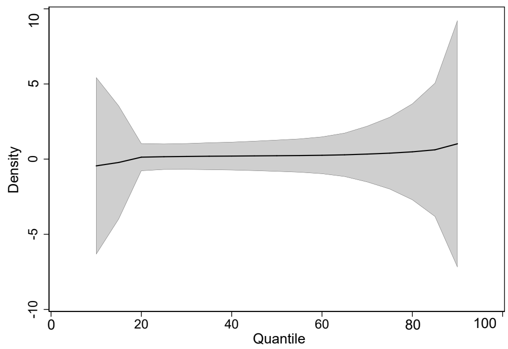

(F) Import intensity

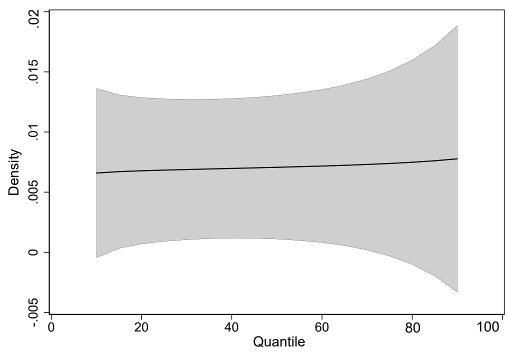

(G) EBIDTA on assets

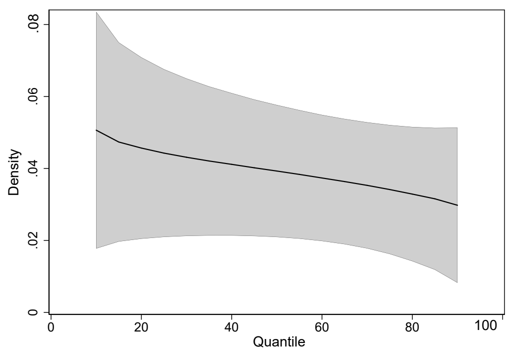

(H) Gross operating surplus on GVA

FIGURE 9: Quantile regressions

Note: The shaded area corresponds to the 95% condidence interval. EBITDA stands for earninsg before interest, taxes, depreciations and amortizations. In order to correct for extreme observations, all dependent variables were subject to a winsorization procedure affecting observations up to the percentile 5 and above percentile 95, except in the cases of TFP and labour productivity, where the percentiles were 1 and 99.

6. Final remarks

Robots and automation in general are an important dimension of the ongoing digital transition. Impacts extend from the labour market to the productive structure and international trade, as changes in the pattern of comparative advantages and firms' location decisions are likely to occur. The analysis of robots' impacts in firms and the overall economy will continue as further data becomes available. Detailed firm-level information on the usage of robots is still sparse, which is a problem when the aim is to establish causal relationships.

This article presents evidence regarding the prevalence of robots in Portuguese firms. The first part of the study frames Portugal in the international context and describes the distribution of robots along sectors and tasks performed. However, the association of robots with firms' performance variables is more relevant. Therefore, in the second part of the article we statistically test the relationship between the existence of robots at the firm and a set of performance variables, namely productivity, wages, international trade and profitability.

We conclude that Portugal has been increasing the number of robots in operation, but it starts from very low levels and other European laggards have evolved much faster. Robots are mostly present in the manufacturing sector, more specifically on metal products and automotive activities. Using the firm-level data, it is possible to corroborate that the manufacturing sector is the one where a larger share of firms reports the utilization of robots, not only in terms of number of firms but also in terms of turnover and employment. As for the association between the existence of robots at the firm and performance, empirical evidence for Portuguese firms points to a positive relationship with productivity, export intensity and profitability, while the correlation with the labour share is negative. The impact on wages is not statistically different from zero.

The decision to use of robots in the production process is taken by the firms within the setup of a profit maximization problem. There is room for public policy in the promotion of these technologies, notably if there are positive spillovers or barriers to investment. Nevertheless, although there are seemingly positive results from the operation of robots in firms, the success of these investments ultimately lies on firms' abilities and business plans. Therefore, subsidization must be carefully designed and the removal of framework costs and barriers to financing and investment should be prioritized.

References

Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo (2020). "Competing with Robots: Firm-Level Evidence from France." *AEA Papers and Proceedings*, 110, 383–88.

Acemoglu, Daron and Pascual Restrepo (2018). "Artificial Intelligence, Automation and Work." NBER Working Papers 24196, National Bureau of Economic Research, Inc.

Acemoglu, Daron and Pascual Restrepo (2020). "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6), 2188–2244.

Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel (2020). "What Are the Labor and Product Market Effects of Automation?: New Evidence from France." Tech. rep., Cato Institute.

Alguacil, Maite, Alessia Lo Turco, and Inmaculada Martínez-Zarzoso (2022). "Robot adoption and export performance: Firm-level evidence from Spain." *Economic Modelling*, 114, 105912.

Almeida, Derick and Tiago Neves Sequeira (2023). "Are Robots, Software, ICT and physical capital related to productivity? A panel quantile approach." *Economics of Innovation and New Technology*, 0(0), 1–18.

Amador, João and Cátia Silva (2023). "A view on ICT and digitalization in Portuguese firms." *Banco de Portugal Economic Studies*, IX(4), 31–52.

Autor, David H. (2015). "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3), 3–30.

Bonfiglioli, Alessandra, Rosario Crinò, Harald Fadinger, and Gino Gancia (2020). "Robot Imports and Firm-Level Outcomes." Discussion Paper Series 2020-243, University of Bonn.

Candeias, Marta, Nuno Boavida, and António Brandão Moniz (2022). "Automation Trends In Portugal: Implications In Productivity And Employment." Estudos de Temas Económicos 165, GEE.

de Vries, Gaaitzen J., Elisabetta Gentile, Sébastien Miroudot, and Konstantin M. Wacker (2020). "The rise of robots and the fall of routine jobs." *Labour Economics*, 66, 101885.

Deng, Liuchun, Steffen Müller, Verena Plümpe, and Jens Stegmaier (2023). "Robots, occupations, and worker age: A production-unit analysis of employment." Tech. rep.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles (2021). "Should Robots Be Taxed?" *The Review of Economic Studies*, 89(1), 279–311.

International Federation of Robotics (2022). "World Robotics." Tech. rep., International Federation of Robotics.

Koch, Michael, Ilya Manuylov, and Marcel Smolka (2021). "Robots and Firms." *The Economic Journal*, 131(638), 2553–2584.

Leone, Fabrizio (2022). "Multinationals, Robots, and the Labor Share." Working Papers ECARES 2022-17, ULB – Universite Libre de Bruxelles.

Levinsohn, James and Amil Petrin (2003). "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies*, 70(2), 317–341.

Machado, José A.F. and J.M.C. Santos Silva (2019). "Quantiles via moments." *Journal of Econometrics*, 213(1), 145–173.

Rios-Avila, Fernando (2020). "MMQREG: Stata module to estimate quantile regressions via Method of Moments." Statistical Software Components, Boston College Department of Economics.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. X, No. 2

Non-technical summary

April 2024

Euro area inflation expectations: A focus on consumers' expectations

Sandra Gomes, Nuno Monteiro and Pedro Pires Ribeiro

Inflation expectations are the rate at which economic agents expect prices to rise in the future. Inflation expectations are central to macroeconomic theory, as they influence critical intertemporal decisions. Moreover, the expectations of financial market participants impact asset prices. Hence, inflation expectations are an important component of the information set guiding monetary policy authorities in pursuing their price stability objective.

While inflation expectation measures may be obtained from various sources, they have been more commonly collected from financial instruments and surveys from professional analysts. To a large extent, this results from data unavailability regarding surveys of other types of agents. Until recently, the primary source for assessing consumers' inflation perceptions and expectations in the euro area was the Business and Consumer Surveys (BCS) conducted by the European Commission (EC) since 2004. However, the European Central Bank (ECB) introduced the Consumer Expectations Survey (CES) in January 2020, gathering monthly data from several euro area countries. This survey is a relevant step to fill knowledge gaps regarding euro area household sector analysis. This article analyses the results of the CES regarding inflation perceptions and expectations. While prior analyses have already focused on the CES, this article covers the period up to December 2023, thereby including the recent surge in inflation and subsequent fall. Despite the richness of the survey, it is noteworthy to highlight that it is relatively recent, so we also provide a detailed comparison with the quantitative inflation perceptions and expectations from the EC's BCS, which is available since 2004. Despite the differences in the design of the two surveys, their comparison is informative, as they focus on the same type of agent.

Figure 1 displays the series for inflation perceptions and 1-year ahead inflation expectations of both surveys and compares with observed inflation. The series exhibit a large co-movement, but compared to observed inflation there seems to be a positive bias, in particular in the case of the BCS.

Using publicly available microdata (i.e., data for responses at the individual level of the respondent) from the CES, we investigate factors influencing consumers' inflation perceptions and expectations. In general, our analyses suggest that consumers in the euro area tend to report an upward bias relative to observed inflation, evident in

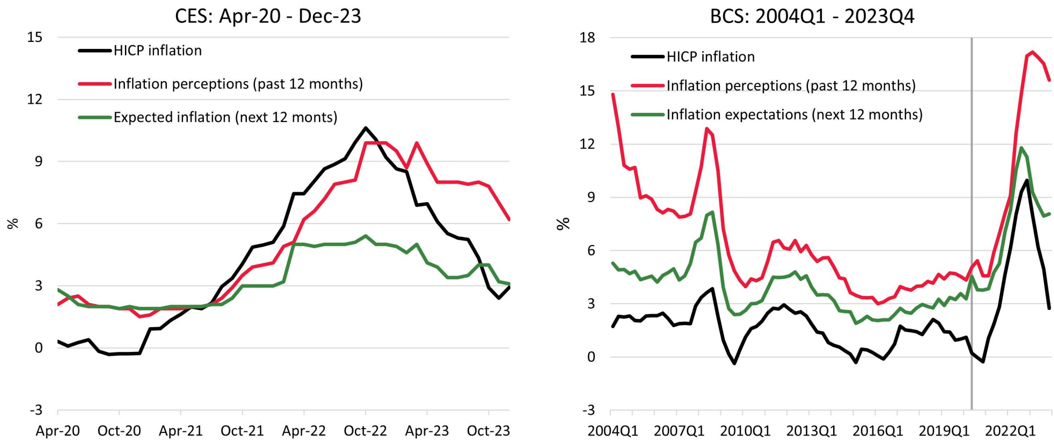

FIGURE 1: Inflation perceptions, 1-year ahead expectations and observed inflation

Sources: ECB, EC, and Eurostat.
Notes: In the right-hand figure, the bar represents the beginning of the CES sample period. The figure shows median inflation perceptions/expectations.

both their perceptions and expectations. In addition, our empirical findings reveal a link between inflation perceptions and observed inflation, and that 1-year inflation expectations are related to past expectations, inflation perceptions and social and demographic characteristics (such as age and income).

Finally, we compute euro area real interest rates by adjusting the nominal interest rate through various measures of 1-year ahead inflation expectations. We then compare these different measures against an equilibrium real rate, as this difference is informative regarding the monetary policy stance. Overall, while alternative measures of the real interest rate based on different inflation expectation measures exhibit similarities, they may lead to different conclusions as for the degree of accommodation or restrictiveness of monetary policy over time.

In summary, we conclude that the CES represents a valuable addition to the array of euro area inflation expectation measures that policymakers should continue to monitor in the future, as it is unfeasible to determine which survey provides the best measure of consumers' inflation expectations.

Euro area inflation expectations: A focus on consumers' expectations

Sandra Gomes
Banco de Portugal
ISEG-ULisboa - UECE, REM

Nuno Monteiro
Banco de Portugal

Pedro Pires Ribeiro
Banco de Portugal
ISCTE-BRU

April 2024

Abstract
Central banks routinely analyse measures of inflation expectations to gain insights into the private sector's perspectives on inflation developments and to evaluate the credibility of monetary policy in fulfilling its mandate. While the analysis of inflation expectations typically focuses on market-based measures and professional forecasters' surveys, the expectations of consumers and firms are also considered relevant as they are believed to influence economic decisions. This article examines inflation expectations among euro area consumers, focusing on the European Central Bank's Consumers Expectations Survey and comparing it with the European Commission's Business and Consumers Surveys. In line with previous results in the literature, the empirical findings confirm an upward bias in consumer's inflation perceptions and expectations relative to actual inflation, and support the link between observed inflation, perceptions, and 1-year ahead expectations. Furthermore, consumers' inflation expectations differ according to age and income. Finally, the article illustrates the use of diverse inflation expectation measures, including those derived from consumers, in computing real rates, emphasising their utility in assessing the monetary policy stance. (JEL: D12, D84, E31, E52, H31)

1. Introduction

Inflation expectations are the rate at which agents, including consumers, firms, financial market analysts, and investors, expect prices of goods and services to change in the future. Inflation expectations are central to macroeconomic theory. Optimal intertemporal decisions (including price and wage setting and consumption-saving decisions) hinge on real variables, i.e., nominal variables adjusted

Acknowledgements: Sandra Gomes ackowledges REM financial support of FCT (Fundação para a Ciência e a Tecnologia), project UIDB/05069/2020. Pedro Pires Ribeiro acknowledges support of FCT, project UIDB/00315/2020. The authors thank the editor (Pedro Duarte Neves), Nuno Alves, Alexandre Carvalho, Sónia Costa, João Quelhas, Paulo Rodrigues, and António Rua for their comments and suggestions. The analyses, opinions and conclusions expressed herein are the sole responsibility of the authors and do not necessarily reflect the opinions of Banco de Portugal or the Eurosystem.

E-mail: sgomes@bportugal.pt; nmmonteiro@bportugal.pt; ppribeiro@bportugal.pt

for expected inflation (see, e.g., ECB (2011)). In various macroeconomic models, inflation expectations, namely those influencing consumption and investment decisions, are a key determinant of macroeconomic variables, including inflation. Expectations of financial market agents have been shown to affect asset prices, such as stock prices and interest rates (see, e.g., Bernanke and Kuttner (2005)). Additionally, the effectiveness of the central bank's communication and the public's confidence in its monetary policy can be gauged by resorting to measures of inflation expectations (see e.g., Christelis *et al.* (2020) and Mellina and Schmidt (2018)). Hence, inflation expectations are an important component of the information set guiding monetary policy authorities in pursuing their inflation objective, while ensuring that long-run inflation expectations remain well-anchored at that goal.

Inflation expectation measures may be obtained from various sources, including financial instruments and surveys. On the one hand, inflation-linked securities (such as inflation-linked bonds, inflation-linked swaps, and inflation-linked options) serve as important real-time monitoring measures for inflation expectations, though they are prone to include risk premia (e.g., inflation risk premium[1]). Still, their interpretation under the assumption of risk neutrality yields a measure of inflation expectations that corresponds to the average expected inflation rate over the security's term. On the other hand, surveys of different types of agents (e.g., market analysts, consumers, and firms) are typically available at a lower frequency than market data. Surveys of professional analysts provide direct measures of inflation expectations for various time horizons but are released with a monthly or quarterly frequency. There are also surveys that monitor expectations of consumers and firms. In fact, these may be particularly relevant as they reflect the expectations that likely influence these agents' economic decisions and may differ from financial market participants'/professional forecasters' expectations.

Monetary authorities routinely monitor measures of inflation expectations both at short and long horizons. The information set of central banks most frequently includes measures based on financial market instruments and professional analysts. To a large extent, this results from data unavailability regarding surveys of other types of agents, as recognised by Gomes *et al.* (2021), although this information gap has been closing. In addition, it has been argued that consumers may be poorly informed (see, e.g., Coibion and Gorodnichenko (2015)). Their answers tend to be volatile and dispersed, they tend to overestimate observed inflation and they appear to be strongly influenced by their own subjective experience (see, e.g., D'Acunto *et al.* (2021)). Fundamentally, there is still limited knowledge about consumers' inflation expectations.

If one believes there is some role for expectations in decision making, then following consumers' and firms' inflation expectations becomes relevant, namely for policymakers. In addition, the monitoring of inflation expectations is important to assess the stance of monetary policy. In fact, the stance of interest rate policy can be gauged by comparing the level of real interest rates, i.e., nominal rates adjusted by a measure

1. The inflation risk premium corresponds to the compensation investors demand for bearing risks related to the uncertainty surrounding future inflation.

of inflation expectations, with that of the equilibrium real rate often called the natural interest rate. The availability of diverse measures for inflation expectations, reflecting the perspectives of different agents, enables the computation of distinct real rates, thereby providing hints regarding different perceptions of the monetary policy stance.

Up to recently, the only source of consumer's inflation expectations measures for the euro area was the Business and Consumer Surveys (BCS) conducted by the European Commission (EC) at quarterly frequency. However, in January 2020, the European Central Bank (ECB) launched a pilot Consumer Expectations Survey (CES) which collected information at a monthly frequency on the perceptions and expectations of euro area consumers across several economic dimensions, including inflation. The pilot phase confirmed the overall high quality of the CES data and ended in June 2021, when a new development phase with further enhancements started, such as an increased sample size, a broader country and topical coverage, among other measures aimed at strengthening the quality of CES analysis at both the aggregate and micro levels. This survey represents a relevant step to fill knowledge gaps regarding euro area household sector analysis.

The contribution of this study is threefold. Firstly, it analyses the results of the CES regarding inflation perceptions and expectations. While prior analyses have already focused on the CES (e.g., Bańkowska *et al.* (2021)), this article covers the period up to December 2023, thereby including the recent surge in inflation and subsequent fall. Despite the richness of the survey, it is noteworthy to acknowledge it is relatively new, so this article also provides a comparison with the quantitative inflation perceptions and expectations from the EC's BCS, which is available since 2004. Even though there are differences between the design of the two surveys, it is informative to compare them as they give information regarding the same type of agents. Secondly, using publicly available microdata (i.e., data at the individual level) from the CES, this article explores some variables affecting consumes' inflation perceptions and expectations at the individual level. Thirdly, acknowledging the pivotal role of inflation expectations, this article computes several euro area real interest rate gaps based on various measures of 1-year ahead inflation expectations and provides an interpretation in terms of the monetary policy stance.

This article is structured as follows. Section 2 examines the euro area consumers' inflation expectations derived from the CES and the BCS. Section 3 analyses the microdata from the CES, aiming to illustrate factors that influence individual inflation perceptions and expectations. In Section 4, different measures of inflation expectations are compared in terms of implications for the monetary policy stance, in particular through the calculation of real interest rates and comparison to the euro area natural rate. Finally, Section 5 presents some concluding remarks.

2. Consumers' inflation expectations in the euro area

The methodologies of opinion surveys, in particular surveys of consumers, exhibit some differences, for example regarding sample representativeness, the way information is collected and how questions are formulated. In addition, survey results are influenced by several factors, including individual characteristics (such as age or income), potential overemphasis on specific personal experiences, and the experience in survey participation.[2]

The next two sub-sections offer a concise overview of the ECB's CES and the EC's BCS and compare the information on consumers' inflation expectations of the two surveys.

2.1. Characterisation of the surveys

The ECB's CES is an online survey initiated in 2020, conducted monthly by the ECB (implementation was outsourced to Ipsos Public Affairs), encompassing both monthly and quarterly questions.[3] The target population of the CES is the population aged 18 and above, residing in the countries included in the CES sample. The sample aims to be representative by age, gender, and region.[4] During the initial (pilot) phase, the survey targeted approximately 10,000 respondents, and this increased to around 14,000 as of end-2023. The euro area aggregates are computed based on existing country coverage, which, up do December 2023, included Belgium, Germany, Spain, France, Italy, and the Netherlands.[5] The main aggregate results of the CES are published monthly on the ECB's website whereas microdata is published on the website following a quarterly release schedule.

This survey includes, on a monthly basis, quantitative questions regarding consumer perceptions about past inflation as well as expectations about future inflation, namely:

How much higher (lower) do you think prices in general are now compared with 12 months ago in the country you currently live in? Please give your best guess of the change in percentage terms. You can provide a number up to one decimal place. ___._%

How much higher (lower) do you think prices in general will be 12 months from now in the country you currently live in? Please give your best guess of the change in percentage terms. You can provide a number up to one decimal place. ___._%

These questions refer to changes in prices in general instead of using the term "inflation" (or "deflation") to avoid the need of familiarity with these economic concepts. Besides inflation, which is the focus of this article, the survey also covers other themes, namely income and consumption, labour markets and economic growth, and housing markets and credit. Regarding inflation, the survey covers information on

2. Repeating survey participants tend to exhibit lower inflation expectations and reduced uncertainty compared to new participants, which suggests that repeating participants learn and become more informed about inflation as they engage more with the survey (see, e.g., Kim and Binder (2023)).

3. The fact that the survey is conducted online may lead to a bias towards more educated people. For a description of the ECB's CES see Bańkowska *et al.* (2021) and Georgarakos and Kenny (2022).

4. While respondents can be aged 70 or above, the requirements for sample representativeness were initially set to include only the 18-70 age range, given the difficulty of recruiting participants aged 70+.

5. From January 2024 onwards, the euro area aggregate includes Austria, Ireland, Finland, Greece, and Portugal and these countries also have publicly available microdata.

inflation perceptions over the past 12 months and inflation expectation for the next 12 months, as well as 3 years ahead. Public data are categorised by: (i) country, (ii) age group and (iii) income quintile.

The EC's BCS is published monthly by the Directorate-General for Economic and Financial Affairs and is derived from surveys conducted by national institutes in the European Union Member States and the candidate countries, based on a common methodology and timetable.[6] The target population of the BCS includes individuals aged 16 and above. The sample size varies across countries according to the heterogeneity of their economies and is generally positively related to their respective population size. In what regards the consumers, around 25,000 individuals are surveyed across the euro area. The aggregate results of the BCS are monthly, published on the EC's website. Some parts of the survey, including consumers' inflation perceptions and expectations, are only available quarterly.

The BCS collects, on a quarterly basis, quantitative questions on inflation perceptions and inflation expectations since 2004 (regularly published since 2019), namely:[7]

By how many percent do you think that consumer prices have gone up/down over the past 12 months? (Please give a single figure estimate): consumer prices have increased by ___._% / decreased by ___._%.

By how many percent do you expect consumer prices to go up/down in the next 12 months? (Please give a single figure estimate): Consumer prices will increase by ___._% /decrease by ___._%.

The survey questions are intentionally vague as regards the meaning of consumer prices, so respondents make their own interpretation as to what basket of goods to consider. In addition to overall figures, results are further disaggregated for specific categories, including by: (i) income level, (ii) gender, (iii) age group, (iv) occupation, and (v) education level.

2.2. *Comparison between surveys*

This sub-section compares the information on consumers' inflation expectations of the two surveys. The analysis focuses on inflation perceptions and 1-year ahead inflation expectations, as they are available in the two surveys. In general, the analysis of the surveys is mainly focused on median figures (rather than the mean) as they are less sensitive to outliers and this choice follows common practice in this literature.

It is a well-established fact in the literature that consumer surveys' results tend to present a bias relative to observed inflation[8] (see, e.g., Arioli *et al.* (2017), Abildgren and Kuchler (2021) and Weber *et al.* (2022)). The literature relates an upward bias of

6. For a description of the EC's BCS methodology see EC (2024).

7. The EC's BCS also includes qualitative questions about inflation developments over the last and the next 12 months, which were also explored in previous literature (see, e.g., Berk (2002) and Dias *et al.* (2010)).

8. Inflation is defined as the year-on-year change of the Harmonised Index of Consumer Prices (HICP).

consumers surveys on inflation to various factors (see, e.g., Coibion and Gorodnichenko (2015) and Trehan (2011)). For example, consumers tend to pay more attention to their own cost of living/shopping basket and their experience regarding observed inflation over their life, or households tend to be more inclined to give greater attention to items that they purchase frequently or that have a large weight in their consumption basket, such as food and gasoline (see, e.g., Weber *et al.* (2022)). Additionally, individuals have a downward-biased memory of the past level of prices, i.e., thinking that prices were lower in the past than what they actually were (see, e.g., Weber *et al.* (2022)).

Table 1 shows that, on average, inflation perceptions are higher than observed euro area HICP inflation in the BCS and the CES, with a more pronounced difference in the former. It is worth noting that, despite posing a similar question to the same type of respondents, the BCS responses concerning both inflation perceptions and expectations are considerably higher than those reported in the CES, particularly during the overlapping period. This discrepancy could be attributed to the design differences between the two surveys, such as the CES being conducted online, which provides easier access to information, and its use of recurring respondents who may gain experience and refine their responses over time, potentially aligning them more closely with observed inflation. While the results of a survey are an estimate of ”true” consumer inflation expectations, the proximity of CES responses to the actual measured levels of inflation, such as the HICP, highlights its potential usefulness in enriching the set of available household’s inflation views. Nevertheless, as can be seen in Figure 1, the existing co-movement between inflation perceptions and observed inflation is also an indication of its usefulness.

		BCS		CES
		Since 2004Q1	*Since 2020Q2*	*Since April 2020*
Inflation perceptions past 12m	Mean (%)	9.8	14.4	6.4
	Median (%)	7.1	10.7	5.2
	P75-P25 (p.p.)	8.1	11.3	6.4
Inflation expectations next 12m	Mean (%)	6.5	10.2	4.6
	Median (%)	4.4	7.3	3.3
	P75-P25 (p.p.)	5.6	8.9	5.9
Observed Inflation	Mean (%)	2.1	4.4	4.4

TABLE 1. CES and BCS descriptive statistics for the euro area

Sources: ECB, EC, and authors’ calculations.
Notes: Mean/Median across consumers’ replies. Interquartile range – difference between the 75 and 25 percentiles of the distribution of inflation perceptions/expectations.

Figure 1 illustrates that while observed inflation is consistently below inflation perceptions in the BCS, in the CES this is only the case from April 2020 to April 2021 and afterwards this bias inverted. In fact, only in April 2021, when HICP inflation had already risen to 2%, did inflation perceptions and expectations start to rise. Inflation expectations stopped increasing in early 2022, while inflation and inflation perceptions continued to increase further. This behaviour suggests that consumers expected inflation to peak and start declining within the following 12 months. Even though observed

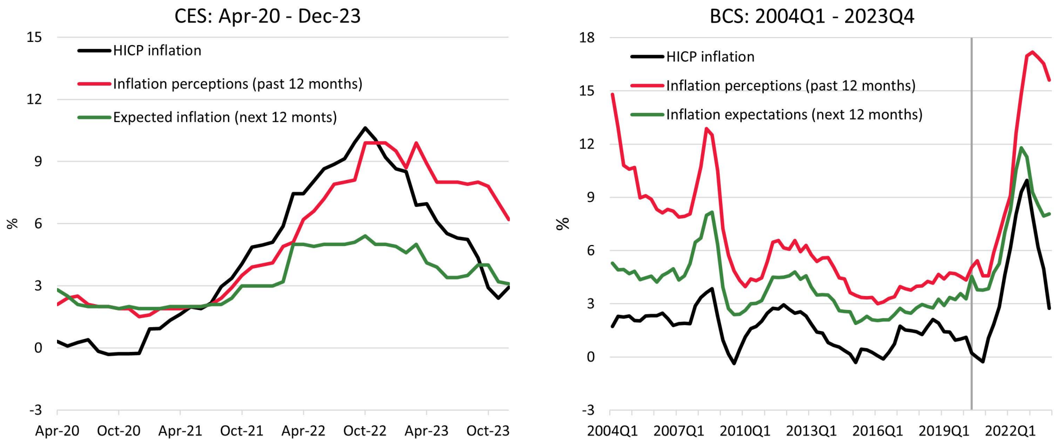

FIGURE 1: Inflation perceptions, 1-year ahead expectations and observed inflation

Sources: ECB, EC, and Eurostat.
Notes: Vertical bar on the right-hand chart identifies the beginning of the CES sample period. Inflation perceptions/expectations – median. The 2nd and 98th percentiles of the weighted distribution of responses for each CES survey round and each country were discarded.

inflation is always lower than BCS inflation perceptions and expectations, the gap widened significantly since mid-2022. The delayed response of inflation perceptions and expectations compared to the recent decline in observed inflation seen in the CES is also noticeable in the BCS. Thus, in both surveys, consumers reported a turn in inflation perceptions and expectations only after the peak in actual inflation, and until December 2023 the decline fell short of the magnitude of the fall of observed inflation.

Table 1 also displays the interquartile range (IQR) of the distribution of both inflation perceptions and expectations which serves as a disagreement measure among consumers about past and future inflation, respectively. The IQR is higher for perceptions compared to expectations and larger in the BCS for the overlapping period with the CES. Figure 2 shows this measure over time. In the BCS there is a clear relationship of IQR with the level of inflation or, equivalently, with the level of perceptions/expectations. This relationship is less clear in the case of the CES, given the relatively short time span. Still, IQR in the CES does seem to have risen in the period when inflation was increasing, especially from September/October 2021 when inflation was around 4% and median inflation perceptions/expectations were already displaying an upward trajectory. In both surveys, IQR remained elevated even after inflation started to fall.

Even though euro area consumers' inflation perceptions seem biased compared to observed inflation, both surveys exhibit a high contemporaneous correlation between inflation perceptions and observed inflation, as shown in Table 2. Also, inflation perceptions and 1-year ahead inflation expectations are highly correlated, as shown in Figure 1.

Previous research has identified substantial heterogeneity in inflation expectations, both across countries and individuals. This variation is related to sociodemographic characteristics such as age, gender, education, and income (see, e.g., Abildgren and

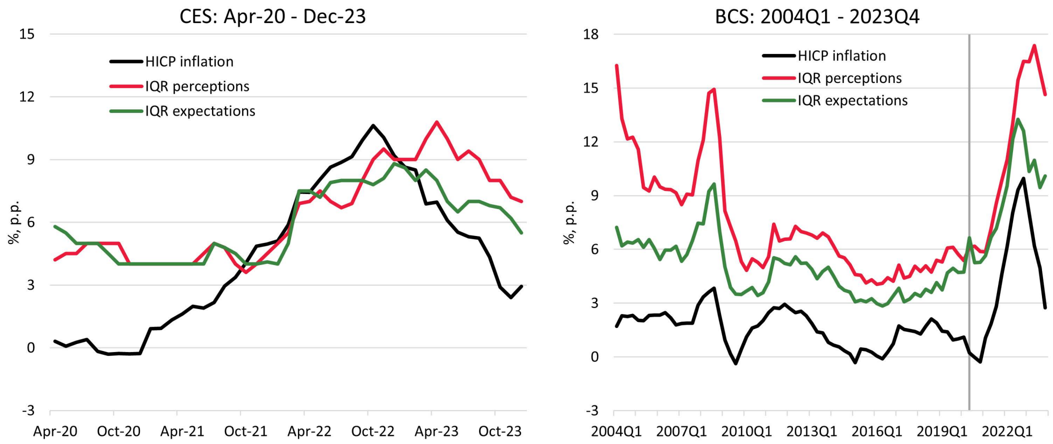

FIGURE 2: Interquartile range in inflation perceptions and 1-year ahead inflation expectations
Sources: ECB, EC, and authors' calculations.
Notes: Vertical bar on the right-hand chart identifies the beginning of the CES sample period. The 2nd and 98th percentiles of the weighted distribution of responses for each CES survey round and each country were discarded.

		Perception vs. Obs. Inflation	Perception vs. Expectation
CES	Since April 2020	0.86	0.90
BCS	Since 2020Q2	0.80	0.85
	Since 2004Q1	0.77	0.89

TABLE 2. Contemporaneous correlations
Sources: ECB, EC, and authors' calculations.

Kuchler (2021); Arioli *et al.* (2017), Bryan and Venkatu (2001); Ehrmann *et al.* (2017) and Jonung (1981)). Next, this study examines inflation perceptions and expectations across two dimensions available in both surveys: age and income brackets. Figure 3 displays consumers' inflation expectations in the CES and the BCS by income and age groups.

The first panel of Figure 3 shows median inflation expectations per income level in both surveys. Consumers with lower incomes generally tend to anticipate higher inflation compared to respondents with higher incomes in the BCS. Differences across income brackets are quite small in the CES, which could be related with the fact that the survey started recently and that the number of respondents has increased gradually. The literature has presented evidence suggesting that agents with more favourable financial situations tend to anticipate lower inflation rates (see, e.g., del Giovane *et al.* (2009); Ehrmann *et al.* (2017); Galati *et al.* (2021)). This result has been linked to a greater degree of uncertainty surrounding future inflation and a more pessimistic assessment of the economic outlook by lower income respondents.

Shifting the focus to the behaviour of inflation expectations across different age groups, the two surveys present divergent findings. In the EC survey, older individuals anticipate lower inflation in the next year, while the results of the CES signal the opposite, i.e., older consumers have higher inflation expectations. In fact, the literature

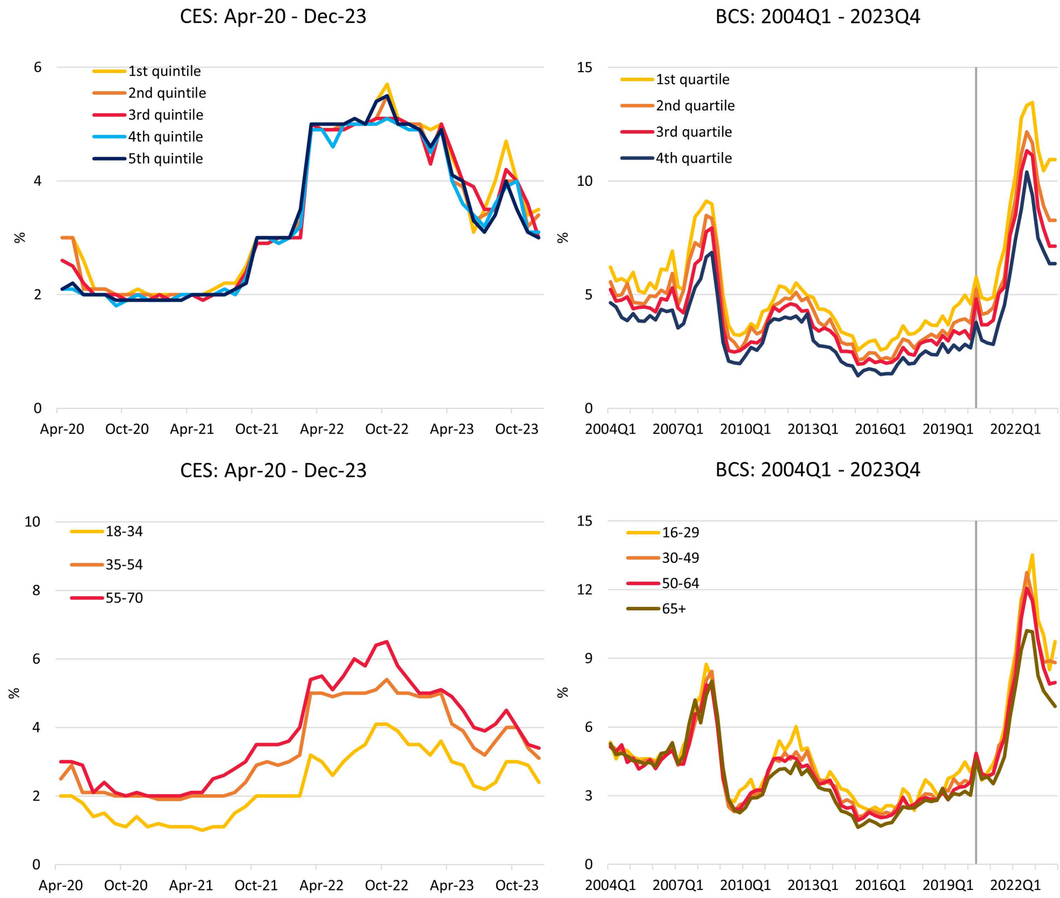

FIGURE 3: Median inflation expectations per income and age brackets

Sources: ECB, and EC.
Notes: Vertical bar on the right-hand chart identifies the beginning of the CES sample period. The figure shows median inflation expectations per income level/age bracket in both surveys. The figure displays publicly available aggregate data for each survey.

on the relationship between age and inflation expectations yields differing results. Some research highlights that inflation expectations are influenced by the inflation history experienced by the respondents (see, e.g., Malmendier and Nagel (2016)). Older respondents base their perceptions on data spanning their entire lifetime, thus covering periods of both low and high inflation, which results in less sensitivity to current inflation rates. In turn, younger consumers' expectations will depend more on the level of inflation seen in the years closer to the date at which the survey is conducted, given their shorter lifetimes. Therefore, if younger respondents predominantly experienced a high inflation period, their perceived inflation expectations will tend to be higher, as Jonung (1981) found in a survey of Swedish consumers conducted in the late 1970s following decades of high inflation. Diamond *et al.* (2020) focus on the post-1995 period in Japan (a period of low inflation, so younger generations would have experienced mostly low inflation) and find that inflation expectations increase with age.[9]

3. Deeper analysis of the ECB's CES

This section shifts the analysis to the individual level, conducting a more in-depth exploration of the CES publicly available microdata, aiming to confirm the results of the previous section and to clarify the relationship between inflation expectations and consumers' social and demographic characteristics. Following the ECB's guide to the computation of CES aggregate statistics, we exclude the 2$^{\text{nd}}$ and 98$^{\text{th}}$ percentiles of the weighted distribution of responses for each survey round and country. A panel dataset is then defined with survey results for the available euro area countries, i.e., Belgium, Germany, Spain, France, Italy, and The Netherlands, comprising the period from April 2020 to December 2023. The analyses focus on the answers to the quantitative questions on both inflation perceptions and 1-year ahead inflation expectations,[10] as well as respondents' details regarding age, income level and country. The dataset can be described as an unbalanced panel (i.e., a sample where entities are observed a different number of times), in which a sizeable share of respondents has only responded sporadically. While this could raise some concerns about the robustness of the empirical exercise, we have replicated the analyses with a restricted sample (considering only respondents who participated in the survey at least 24 editions), and the results are similar to those presented in this section.

3.1. Inflation perceptions and observed inflation

To begin the exploration of the CES microdata, this study analyses the link between inflation perceptions in the CES with actual inflation. Employing a simple linear regression, inflation perceptions are modelled as a function of observed inflation. Formally, this relationship is expressed in equation 1 as follows:

$$\pi_{it}^{P} = \mu_i + \beta\pi_{it} + u_{it} \tag{1}$$

where pi_{it}^{P} corresponds to inflation perceptions of individual i at time t and π_{it} is observed inflation in country of individual i at time t. μ_i is an individual specific effect, which is assumed to be fixed. Current inflation is included as the regressor because the goal is to assess if consumers' perceptions are related to inflation and not whether

9. Juselius and Takats (2015) focus on trend (i.e., low frequency) inflation (not inflation expectations) over the 1955–2010 period for 22 advanced economies and found that the larger the proportion of young and old in the total population, the higher inflation. In other words, when the working-age population is larger, the effect is disinflationary. This result holds for a large number of countries across all time periods.

10. Questions regarding inflation expectations and perceptions ask about prices in the country the respondent lives in. Therefore, when comparing the responses at the individual level to actual inflation outcomes this study considers the respective country inflation.

the available information, namely regarding the HICP, influences their perceptions.[11] The estimation results for the full sample period are detailed in Table 3. The estimate for parameter β in equation 1 is statistically significant at the usual significance levels, indicating that, on average, across the entire sample, a 1 p.p. increase in observed inflation is associated with a 0.55 p.p. increase in perceptions.[12]

Coefficient of obs. inflation (β)	0.55 ***
	(0.01)
N	547850
Within R^2	0.10
Overall R^2	0.08

TABLE 3. Estimation results of equation 1

Notes: This table records the results of the estimation of equation 1 from April 2020 to December 2023. The estimate for the constant is not reported, as it lacks an interpretation when estimating the regression with fixed effects. Robust standard errors in brackets. $*$, $**$ and $***$ denote significance at the 10%, 5% and 1% level, respectively.

Next, considering the substantial changes in actual inflation throughout the full sample as well as the widening gap between actual and perceived inflation since the first quarter of 2022 (see Figure 1), regression 1 is also estimated with an expanding window that starts with data ranging from April 2020 to April 2022 and concludes with data from April 2020 to December 2023, as well as with a rolling 24-month window. Figure 4 depicts the estimation results. The estimates for parameter β increase up to early-2023 and decline thereafter, noticeably in the rolling window regression. In both cases, the estimated coefficient is always positive and statistically significant, confirming that there is a relationship between respondents' inflation perceptions and actual inflation, although varying over time.

3.2. *Factors influencing 1-year inflation expectations*

Next, this study explores in more detail 1-year ahead inflation expectations and investigates not only the link with inflation perceptions, but also whether consumers' social and demographic characteristics play a role in individual inflation expectations. To achieve this, alternative specifications are considered:

$$\pi_{it}^{E} = \mu_i + \omega\pi_{it}^{P} + \vartheta\pi_{it-1}^{E} + u_{it} \;\; i = 1, ..., N \text{ and } t = 1, ..., T \tag{2}$$

$$\pi_{it}^{E} = \alpha + \omega\pi_{it}^{P} + \vartheta\pi_{it-1}^{E} + \beta X_{it} + u_{it} \;\; i = 1, ..., N \text{ and } t = 1, ..., T \tag{3}$$

11. If the goal was to analyse the use of available HICP by consumers to form expectations, then it would make sense to include lagged inflation as a regressor, because at the time of the survey for a given month the corresponding HICP data has not been released.

12. Equation 1 was also estimated including the 1-month lagged inflation perception as a regressor and the estimates for all the parameters are statistically significant at the conventional levels.

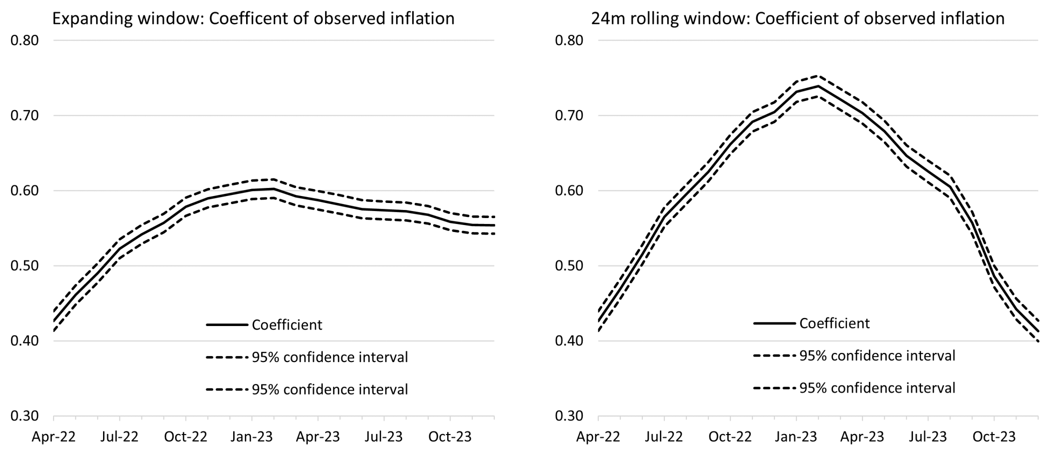

FIGURE 4: Estimation results for equation 1 with an expanding window and a rolling window

where π_{it}^{E} corresponds to consumers' inflation expectations 12-months ahead at time t for individual i, πit^{P} corresponds to consumers' inflation perceptions in the last 12 months at time t for individual i, N and T are the number of cross sections and time periods, respectively. In equation 2, μ_i is an individual specific effect, which is assumed to be fixed. In equation 3, fixed effects are discarded, and a row vector of exogenous explanatory variables related to social and demographic features of dimension k denoted by X_{it} is added, with β representing a k by 1 vector of coefficients.

Given the potential link between inflation perceptions and 1-year ahead inflation expectations, equation 2 is firstly estimated by means of a regression where the only regressor are inflation perceptions. Table 4 shows the results of this specification in regression A, confirming that 1-year ahead inflation expectations are positively and significantly correlated with inflation perceptions at the individual level.

Regression B expands the previous regression by adding 1-year ahead inflation expectations lagged by 1 period as a regressor. Inflation perceptions remain statistically significant and the estimate for the parameter of lagged inflation expectations (ϑ) is also statistically different from zero. Thus, this regression points to some persistence in expectations regarding inflation 12-months ahead.

Equation 3 allows to test the role of consumers' social and demographic features, as described in Section 2. First, regression C1 expands regression B with the age group of each consumer. Regression C2 includes information on the income level of the respondent instead of its age, while regression C3 includes both age and income variables. Overall, the results broadly support the view that older consumers and those with lower incomes tend to expect higher inflation expectations. These estimations confirm the divergence found in Section 2.2 between the CES and the BCS regarding the role of age in inflation expectations. The relationship between income level and inflation expectations in the CES can also be interpreted more clearly, pointing to the same inverse relation between income and inflation expectations that was present in the BCS results.

Considering regression C3, we repeat this estimation for the individuals of each country separately. The results are presented in Table 5. Overall, the main conclusions reported for the euro area still apply for each individual country: (i) the estimates for the

	Regressions				
	A	B	C1	C2	C3
Constant term			0.74 ***	1.11 ***	0.93 ***
			(0.03)	(0.03)	(0.04)
Inflation perceptions	0.42 ***	0.39 ***	0.37 ***	0.37 ***	0.37 ***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
1m lagged inf. expectations		0.15 ***	0.23 ***	0.23 ***	0.23 ***
		(0.00)	(0.00)	(0.00)	(0.00)
Age (omitted: 18-34 Years-old)					
35-49 Years old			0.19 ***		0.22 ***
			(0.03)		(0.03)
50-70 Years old			0.26 ***		0.29 ***
			(0.04)		(0.04)
Over 71 Years old			0.38 ***		0.40 ***
			(0.06)		(0.06)
Income quintiles (omitted: 1^{st})					
2^{nd} quintile				-0.10 **	-0.10 **
				(0.04)	(0.04)
3^{rd} quintile				-0.16 ***	-0.17 ***
				(0.04)	(0.04)
4^{th} quintile				-0.31 ***	-0.33 ***
				(0.04)	(0.04)
5^{th} quintile				-0.40 ***	-0.42 ***
				(0.04)	(0.04)
N	547850	446279	446279	446279	446279
Within R^2	0.25	0.27	0.27	0.27	0.27
Overall R^2	0.40	0.48	0.50	0.50	0.50

TABLE 4. Estimation results of equation 2 over the full dataset

Notes: This table records the estimates for specifications 2 and 3 from April 2020 to December 2023 for the euro area. Robust standard errors in brackets. $*$, $**$ and $***$ denote significance at the 10%, 5% and 1% levels, respectively.

parameter associated with inflation perceptions remain significant at conventional levels in all cases; (ii) the estimates for the parameter ϑ are statistically significant at the 1% level, supporting specifications that include some persistence in inflation expectations; (iii) the impact of social and demographic features on the behaviour of 1-year inflation expectations is also relevant at the country level, in particular in the case of the age brackets, but can be less compelling in some countries.

4. 1-year ahead inflation expectations measures and real interest rate gaps

As argued above, inflation expectations influence economic decisions, such as wage and price setting, as well as investment, saving, and consumption choices. Thus, inflation expectations play a key role in the monetary policy transmission mechanism. Economic decisions hinge on the real interest rate, which is the nominal return adjusted for expected inflation (see, e.g., Armantier *et al.* (2015) and Duca-Radu *et al.* (2021)). An expected rise (fall) in inflation in future periods leads to a decline (rise) in the real interest rate. However, there are challenges related to the computation of the real interest rate:

	EA	BE	DE	ES	FR	IT	NL
Constant	0.93 ***	0.61 ***	0.84 ***	1.24 ***	0.91 ***	0.92 ***	0.84 ***
	(0.04)	(0.11)	(0.07)	(0.10)	(0.06)	(0.11)	(0.10)
Inflation perceptions	0.37 ***	0.36 ***	0.43 ***	0.34 ***	0.45 ***	0.36 ***	0.34 ***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
1m lagged infl. expect.	0.23 ***	0.25 ***	0.19 ***	0.21 ***	0.18 ***	0.26 ***	0.24 ***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Age (omitted: 18-34 years)							
35-49 years	0.22 ***	0.22 **	0.08	0.28 ***	0.06	0.49 ***	0.19 **
	(0.03)	(0.09)	(0.05)	(0.09)	(0.06)	(0.10)	(0.09)
50-70 years	0.29 ***	0.37 ***	0.14 **	0.33 ***	0.15 **	0.55 ***	0.26 ***
	(0.04)	(0.10)	(0.06)	(0.10)	(0.06)	(0.12)	(0.09)
Over 71 years	0.40 ***	0.34 *	0.20 **	0.75 ***	0.13	0.83 ***	0.38 ***
	(0.06)	(0.17)	(0.09)	(0.20)	(0.09)	(0.19)	(0.15)
Inc. quint. (omitted: 1st)							
2nd quintile	-0.10 **	0.13	0.03	-0.29 ***	-0.24 ***	-0.03	0.05
	(0.04)	(0.12)	(0.07)	(0.10)	(0.07)	(0.13)	(0.12)
3rd quintile	-0.17 ***	-0.15	0.01	-0.39 ***	-0.23 ***	-0.01	-0.14
	(0.04)	(0.12)	(0.07)	(0.11)	(0.07)	(0.13)	(0.11)
4th quintile	-0.33 ***	-0.20 *	-0.13 *	-0.30 ***	-0.30 ***	-0.76 ***	-0.15
	(0.04)	(0.11)	(0.07)	(0.11)	(0.07)	(0.13)	(0.11)
5th quintile	-0.42 ***	-0.27 **	-0.21 ***	-0.70 ***	-0.30 ***	-0.65 ***	-0.04
	(0.04)	(0.11)	(0.07)	(0.11)	(0.07)	(0.12)	(0.11)
N	446279	34376	91011	93247	94113	99639	33893
Within R^2	0.27	0.31	0.33	0.22	0.29	0.27	0.32
Overall R^2	0.50	0.54	0.50	0.45	0.49	0.50	0.49

TABLE 5. Estimation results of regression C3 for each country

Notes: This table records the estimates for equation 3 from April 2020 to December 2023, by country. Robust standard errors in brackets.$*$, $**$ and $***$ denote significance at the 10%, 5% and 1% level, respectively. EA – Euro area; BE – Belgium, DE – Germany, ES – Spain, FR – France, IT – Italy, NL – The Netherlands.

while the nominal interest rate choice usually focuses on the short-term horizon (e.g., the policy rate or a short-term market rate), the selected measure of inflation expectations is not straightforward. In fact, the availability of diverse measures of inflation expectations, for example representing different agents' perspectives, allows for the computation of distinct real interest rates.

In general, the primary monetary policy instrument is the policy rate (see, e.g., ECB (2021)). By steering the level of (very) short-term interest rates, monetary policy influences the level of other nominal interest rates and real interest rates. If the observed real interest rate is above the equilibrium real interest rate, frequently called the natural rate,[13] then policy is said to be restrictive. Conversely, if it is below the natural rate, i.e., there is a negative real interest rate gap, then policy is deemed accommodative. However, employing the natural real interest rate as a benchmark

13. Various definitions of the natural interest rate are present in the literature. The foundational concept was introduced by Wicksell (1898) and can be characterised as (i) the interest rate that ensures the equilibrium in the savings/investment markets, (ii) the rate associated with the return on capital, and (iii) the interest rate, which is neutral in respect to prices, as tends neither to raise or lower them. In general, the natural interest rate tends to be interpreted as the real interest rate that would prevail under circumstances considered as desirable from the standpoint of macroeconomic stabilisation.

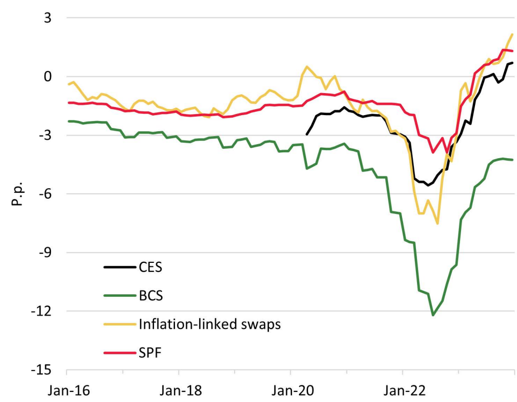

FIGURE 5: Euro area real interest rate gaps

Sources: ECB, EC, LSEG and authors' calculations.
Notes: Real interest rate gaps – differential between real interest rates, (computed as the difference between a nominal interest rate and inflation expectations) and the natural interest rate (median of estimates reported in related box of Brand *et al.* (2024)). Nominal interest rate: €STR from Oct-19 onwards, backdated using the EONIA rate.

for assessing monetary policy stance presents at least two major challenges: it is an unobservable variable, and its level evolves over time. Inferences about the natural interest rate are subject to high uncertainty, as its estimates and their interpretation are model and data dependent. As there are alternative ways to estimate the natural interest rate, this study considers a range of measures recently published by the ECB (Brand *et al.* (2024)) and uses the median of those estimates. It is also worth highlighting that some of the natural rate measures provided by the ECB are based on macroeconomic models, which contain measures of inflation expectations that are not necessarily the same as the ones considered here. This exercise acknowledges but does not tackle the uncertainty surrounding the estimation of the natural rate. Instead, this study focuses on the implications of considering different inflation expectations to the assessment of the euro area monetary policy stance.

Several real interest rate gaps for the euro area are calculated as the difference between different real interest rates and the natural rate. In practice, the differences among these measures result solely from the different measures of inflation expectations: those implied by the two consumer surveys (BCS and CES), financial market instruments (inflation-linked swaps), and the ECB Survey of Professional Forecasters (SPF).[14] Each of these measures can be seen as the perceived stance of monetary policy by each group of economic agents. Figure 5 displays the various alternative measures of real interest rate gaps in the euro area.

The use of different inflation expectation measures yields somewhat different real interest rate gaps. Since 2016, all available measures show that monetary policy stance

14. The SPF collects information on expected inflation rates, real GDP growth, and unemployment levels across various time-frames within the euro area, presenting both point forecasts and probability distributions to quantify risk and uncertainty, with publication occurring quarterly since 1999.

was consistently expansionary, though as expected the degree varies. In the pandemic crisis period, all measures show a drop, but this is considerably stronger in the BCS. Since mid-2022, all measures show an increase in the real rate gap, that only in the case of the BCS remains negative. Given the unobservable nature of consumers' expectations, it is unfeasible to determine which survey provides the best measure of consumer inflation expectations. Therefore, the CES emerges as a valuable addition to the array of euro area inflation expectation measures that policymakers should monitor.

5. Conclusion

Inflation expectations hold significant importance in so far as they influence the decisions of economic agents. Accordingly, central banks routinely monitor inflation expectation measures to gain valuable insights into the private sector's perspectives on the inflation outlook and to evaluate the credibility of monetary policy. In this vein, central banks typically examine various measures of inflation expectations, encompassing both market-based measures and survey-based measures provided by professional analysts. However, the relevance of inflation expectations extends beyond market participants and professional forecasters, including expectations of households and firms.

This article examines inflation expectations among euro area consumers, with a specific focus on the ECB's CES. In particular, it compares key findings with those of the EC's BCS. The analyses confirm that consumers in the euro area tend to show an upward bias relative to observed inflation, both in terms of inflation perceptions and inflation expectations, thereby corroborating previous research.

The empirical exercises using CES microdata reveal a link between inflation perceptions and observed inflation, and that 1-year inflation expectations are related to past expectations, inflation perceptions and social and demographic characteristics (such as age and income).

Finally, a comparison among various measures of 1-year ahead inflation expectations is conducted. Recognising that monitoring inflation expectations for monetary policy is crucial not only for anticipating inflation developments but also for assessing the effectiveness of the central bank's communication and ultimately the stance of monetary policy, several measures of the real interest rate are computed. This exercise considers a proxy for the (nominal) monetary policy rate at any given moment and different inflation expectation measures. Although some expected similarities exist, the analysis illustrates that the use of different measures of inflation expectations may lead to different conclusions regarding the degree of accommodation or restrictiveness over time. By comparison with the EC's BCS, the CES measure of inflation expectations aligns more closely with two established indicators of inflation expectations (those of professional analysts and market expectations).

In short, this article finds that the CES adds value to the range of inflation expectation measures in the euro area. Policymakers should continue to monitor it alongside other surveys, as it is difficult to determine which one best represents the “true” consumers’ inflation expectations in the euro area.

References

Abildgren, Kim and Andreas Kuchler (2021). "Revisiting the inflation perception conundrum." *Journal of Macroeconomics*, 67(C).

Arioli, Rodolfo, Colm Bates, Heinz Christian Dieden, Ioana Duca, Roberta Friz, Christian Gayer, Geoff Kenny, Aidan Meyler, and Iskra Pavlova (2017). "EU consumers´ quantitative inflation perceptions and expectations: an evaluation." Occasional paper series, European Central Bank.

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar (2015). "Inflation Expectations And Behavior: Do Survey Respondents Act On Their Beliefs?" *International Economic Review*, 56(2), 505–536.

Bańkowska, Katarzyna, Ana Maria Borlescu, Evangelos Charalambakis, António Dias da Silva, Davide di Laurea, Maarten Dossche, Dimitris Georgarakos, Juha Honkkila, Neale Kennedy, and Kenny (2021). "ECB Consumer Expectations Survey: an overview and first evaluation." Occasional Paper Series 287, European Central Bank.

Berk, Jan Marc (2002). "Consumers' Inflation Expectations And Monetary Policy In Europe." *Contemporary Economic Policy*, 20(2), 122–132.

Bernanke, Ben S. and Kenneth N. Kuttner (2005). "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance*, 60(3), 1221–1257.

Brand, Claus, Noëmie Lisack, and Falk Mazelis (2024). "Estimates of the natural interest rate for the euro area: an update." *Economic Bulletin Boxes*, Issue 1.

Bryan, Michael F. and Guhan Venkatu (2001). "The demographics of inflation opinion surveys." *Economic Commentary*, (October).

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, and Maarten van Rooij (2020). "Trust in the Central Bank and Inflation Expectations." *International Journal of Central Banking*, 16(6), 1–37.

Coibion, Olivier and Yuriy Gorodnichenko (2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8), 2644–2678.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber (2021). "Exposure to Grocery Prices and Inflation Expectations." *Journal of Political Economy*, 129(5), 1615–1639.

del Giovane, Paolo, Silvia Fabiani, and Roberto Sabbatini (2009). "What's Behind ´Inflation Perceptions´ A Survey-Based Analysis of Italian Consumers." *Giornale degli Economisti*, 68(1), 25–52.

Diamond, Jess, Kota Watanabe, and Tsutomu Watanabe (2020). "The Formation of Consumer Inflation Expectations: New Evidence From Japan's Deflation Experience." *International Economic Review*, 61(1), 241–281.

Dias, Francisco, Cláudia Duarte, and António Rua (2010). "Inflation expectations in the euro area: are consumers rational?" *Review of World Economics (Weltwirtschaftliches Archiv)*, 146(3), 591–607.

Duca-Radu, Ioana, Geoff Kenny, and Andreas Reuter (2021). "Inflation expectations, consumption and the lower bound: micro evidence from a large multi-country survey." *Journal of Monetary Economics*, 118(C), 120–134.

EC (2024). "The Joint Harmonised EU Programme of Business and Consumer Surveys - User guide (updated January 2024)." Tech. rep., European Commission.

ECB (2011). "Inflation Expectations in the euro area: a review of recent developments." *Economic Bulletin Articles*, 2011.

ECB (2021). "The ECB's monetary policy strategy statement." URL `https://www.ecb.europa.eu/home/search/review/html/ecb.strategyreview_monpol_strategy_statement.en.html`.

Ehrmann, Michael, Damjan Pfajfar, and Emiliano Santoro (2017). "Consumers' Attitudes and Their Inflation Expectations." *International Journal of Central Banking*, 13(1), 225–259.

Galati, Gabriele, Richhild Moessner, and Maarten van Rooij (2021). "The anchoring of long-term inflation expectations of consumers: insights from a new survey." BIS Working Papers 936, Bank for International Settlements.

Georgarakos, Dimitris and Geoff Kenny (2022). "Household spending and fiscal support during the COVID-19 pandemic: Insights from a new consumer survey." *Journal of Monetary Economics*, 129(S), 1–14.

Gomes, Sandra, Nikolay Iskrev, and Pedro Pires Ribeiro (2021). "Euro area inflation expectations during the COVID-19 pandemic." *Banco de Portugal Economic Studies*, 7(4), 3–24.

Jonung, Lars (1981). "Perceived and Expected Rates of Inflation in Sweden." *American Economic Review*, 71(5), 961–968.

Juselius, Mikael and Elod Takats (2015). "Can demography affect inflation and monetary policy?" BIS Working Papers 485, Bank for International Settlements.

Kim, Gwangmin and Carola Binder (2023). "Learning-through-Survey in Inflation Expectations." *American Economic Journal: Macroeconomics*, 15(2), 254–278.

Malmendier, Ulrike and Stefan Nagel (2016). "Learning from Inflation Experiences." *The Quarterly Journal of Economics*, 131(1), 53–87.

Mellina, Sathya and Tobias Schmidt (2018). "The role of central bank knowledge and trust for the public's inflation expectations." Discussion Papers 32/2018, Deutsche Bundesbank.

Trehan, Bharat (2011). "Household inflation expectations and the price of oil: it's déjà vu all over again." *FRBSF Economic Letter*, May23.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion (2022). "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications." *Journal of Economic Perspectives*, 36(3), 157–184.

Wicksell, Knut (1898). *Interest and Prices: A study of the causes regulating the value of Money*. Macmillan, London.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. X, No. 2

Non-technical summary

April 2024

Human Capital and Entrepreneurship

Attila Gyetvai, Nasir Hossein Dad, Nicholas Kozeniauskas and Eugene Tan

Entrepreneurship is a key driver of business dynamism and economic growth and, as such, is of key policy interest. Informed policy decisions require a thorough understanding of the process of entrepreneurship, especially when it comes to barriers that potential entrepreneurs face. One such barrier is the outside option of entrepreneurship: what entrepreneurs could do and how much they could earn if they returned to paid work. Entrepreneurs accumulate business-specific human capital while foregoing human capital accumulation in paid work; therefore, there is reason to believe that the outside option might evolve with entrepreneurship.

Entrepreneurs' options in the labor market, as for any worker, evolve endogenously according to tenure and to the type of activity developed. On the one hand, entrepreneurs accumulate human capital specific to their business while simultaneously foregoing human capital accumulation in paid work. Thus, their paid work-specific human capital stagnates or even erodes while they run a business, which decreases their outside option of returning. On the other hand, their relative disadvantage might diminish if their entrepreneurial human capital can be transferred to paid work. These two forces together determine the endogenous evolution of entrepreneurs' outside option.

We present empirical evidence suggesting that outside options decrease with entrepreneurial experience. We reach these conclusions by comparing the wage trajectories of entrepreneurs who return to paid work (who we term "return-entrepreneurs") with those who never started a business (who we term "never-entrepreneurs").

We find that return-entrepreneurs suffer large and persistent wage losses relative to never-entrepreneurs upon returning to paid work. The immediate wage loss is 18 percent and stands at 5 percent in the long run. Longer entrepreneurial experience is also associated with more persistent losses. Return-entrepreneurs with 5 and 10 years of experience suffer a 3 and 6 percentage point larger wage loss, respectively, than those with only one year of experience. Finally, we find that the losses are larger for more educated people: Return-entrepreneurs without a high school diploma recover to the

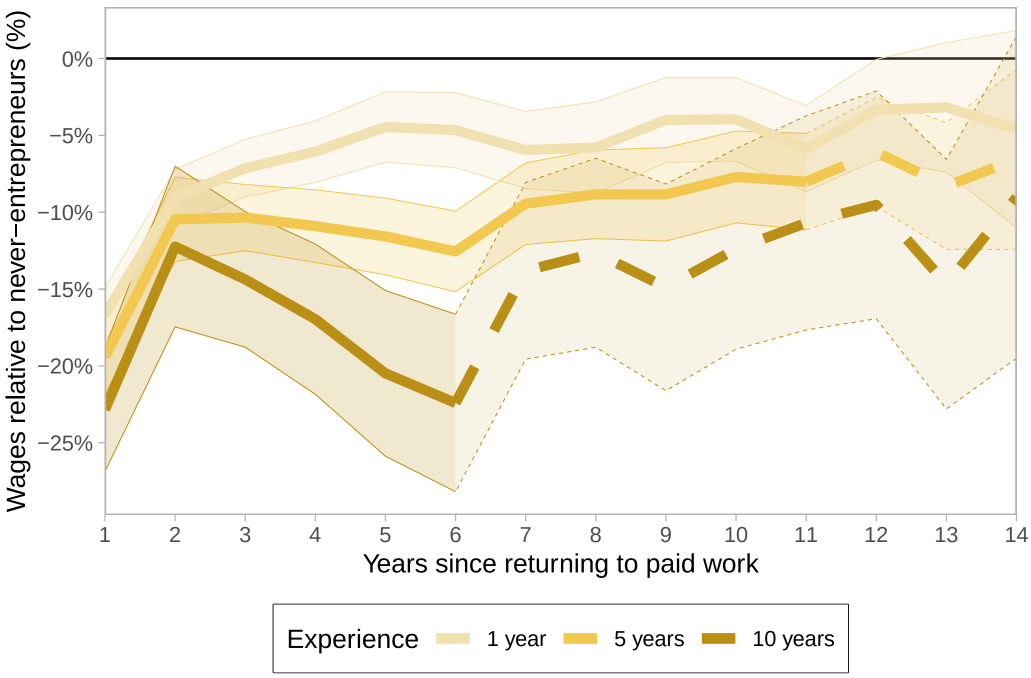

FIGURE 1: Wage trajectories by entrepreneurial experience

Notes: Regression results from specification with occupation FEs. Increasingly darker lines represent wage trajectories after returning to paid work from 1, 5, and 10 years of entrepreneurship, respectively. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

wage level of never-entrepreneurs within 10 years, whereas those with a high school diploma or college degree do not recover within 14 years.

Overall our empirical results suggest that entrepreneurial human capital is specific to business ventures and is not a substitute for human capital gained through experience in paid work.

Banco de Portugal
Banco de Portugal Economic Studies, Vol. X, No. 2

Human Capital and Entrepreneurship

Attila Gyetvai
Banco de Portugal &
IZA

Nasir Hossein Dad
University of Toronto,
Department of Economics

Nicholas Kozeniauskas
Banco de Portugal &
Católica Lisbon SBE

Eugene Tan
University of Toronto,
Rotman School of Management

April 2024

Abstract
We investigate how entrepreneurial human capital shapes the outside option of entrepreneurs when they return to paid work. Using data on Portuguese work histories, we find that entrepreneurs suffer large and persistent wage losses upon returning to paid work: their wage trajectory is 18 percent lower compared to never-entrepreneurs immediately after return and remains 5 percent lower for 14 years. We also show that wage losses increase with entrepreneurial experience, and that return-entrepreneurs with higher levels of education suffer larger and more persistent wage losses. Our results imply that the outside option of entrepreneurs comoves with business-specific entrepreneurial human capital. (JEL: E22, J24, J31, L26)

1. Introduction

Entrepreneurship is a key driver of business dynamism and economic growth and, as such, is of key policy interest. Informed policy decisions require a thorough understanding of the process of entrepreneurship, especially when it comes to barriers that potential entrepreneurs face. One such barrier is the outside option of entrepreneurship: what entrepreneurs could do and how much they could earn if they returned to paid work. Entrepreneurs accumulate business-specific human capital while foregoing human capital accumulation in paid work; therefore, there is reason to believe that the outside option might evolve with entrepreneurship.

In this paper, we present empirical evidence suggesting that outside options decrease with entrepreneurial experience. We reach these conclusions by comparing the wage trajectories of entrepreneurs who return to paid work (who we term "return-entrepreneurs") with those who never started a business (who we term

Acknowledgements: We thank Daniel Luzio Castelo for excellent research assistance. Further thanks to Statistics Portugal for data access and to the Bank of Portugal Microdata Research Laboratory for computing resources. The views expressed here are of the authors and do not necessarily reflect those of the Bank of Portugal or the Eurosystem.

E-mail: attila.gyetvai@bportugal.pt; nasir.hossein@mail.utoronto.ca; nkozeniauskas@bportugal.pt; eugene.tan@rotman.utoronto.ca

"never-entrepreneurs"). Our comparisons rely on return-entrepreneurs and never-entrepreneurs being otherwise similar in terms of their demographic and employment characteristics. We find that return-entrepreneurs suffer large and persistent wage losses relative to never-entrepreneurs upon returning to paid work, and that longer entrepreneurial experience is associated with more persistent losses. Furthermore, we show that return-entrepreneurs with higher levels of education suffer larger and more persistent wage losses. The wage losses of return-entrepreneurs with a high school diploma or college degree worsen with entrepreneurial experience, while they do not vary with experience for those without a high school diploma. Our empirical results suggest that entrepreneurial human capital is specific to business ventures and is not a substitute for human capital gained through experience in paid work.

Our analysis is motivated by the lack of prior evidence on the endogenous evolution of entrepreneurial outside options. The extant literature typically assumes that outside options are constant over the course of entrepreneurship, or evolve independently of tenure in entrepreneurship (Hopenhayn 1992; Melitz 2003; Cagetti and De Nardi 2006; Buera and Shin 2013). However, there is no reason—other than for simplicity and lack of documented evidence—to assume that outside options do not evolve endogenously with entrepreneurial experience. Theoretically, the effect of entrepreneurial experience on outside options could go either way. On the one hand, entrepreneurs accumulate human capital specific to their business while simultaneously foregoing human capital accumulation in paid work. Thus, their paid work-specific human capital stagnates or even erodes while they run a business, which decreases their return option. On the other hand, their relative disadvantage might diminish if their entrepreneurial human capital can be transferred to paid work. These two forces together determine the endogenous evolution of entrepreneurs' outside option.

We provide suggestive evidence on the endogenous evolution of entrepreneurial outside option by estimating the impact of entrepreneurial human capital on labor market outcomes upon returning to paid work. We compare the wage trajectories of return-entrepreneurs to never-entrepreneurs who are otherwise similar. We start by making these comparisons among workers of the same gender, age, and education level in the same calendar year who are working in the same occupation. Next, we compare workers in the same gender, age, education, calendar year, occupation, and sector. Finally, we compare workers in the same gender, age, education, calendar year, occupation, sector, and location. By notching up the tightness of our comparisons, we examine the selection patterns of return-entrepreneurs across sectors and locations. Our analysis exploits rich Portuguese microdata covering the universe of worker histories linked to private firms and firm balance sheets in 2004–2020: these data allow us to track individuals across entrepreneurship and paid work, all the while observing their demographic and employment characteristics.

Taking our empirical approach to the Portuguese data, we find that return-entrepreneurs suffer an 18 percent wage loss relative to their never-entrepreneur counterparts immediately upon returning to paid work. The wage trajectory of return-entrepreneurs remains 5 percent lower than for never-entrepreneurs for 14 years. These results come from comparisons of return-entrepreneurs with never-entrepreneurs in the same occupation. Comparing return-entrepreneurs with never-entrepreneurs in the same occupation and sector, the immediate wage loss is 15 percent and the long-run wage loss is 3 percent: this lower difference in relative terms implies that return-entrepreneurs select into sectors with lower than average wages, thereby recovering faster to a lower baseline. We detect similar composition effects for the location of return-entrepreneurs.

Going further, we show that the wage losses for return-entrepreneurs with longer entrepreneurial experience are larger. Return-entrepreneurs with 1 year of experience suffer a long-run wage loss of 5 percent while it is 3 percentage points larger for those with 5 years of experience and 6 percentage points larger for those with 10 years of experience. Return-entrepreneurs do not catch up to, nor surpass never-entrepreneurs within a time frame that our empirical strategy could detect.

Finally, we repeat our analysis on subsamples by education to understand whether these patterns are stronger for more educated return-entrepreneurs: they are more important from an aggregate perspective as they tend to run larger businesses (Queiró 2022). We find that return-entrepreneurs without a high school diploma experience an immediate wage loss of 13 percent which persists for 10 years, whereas those with a high school diploma or college degree suffer an 18–22 percent immediate loss with no recovery within 14 years. Furthermore, more highly educated return-entrepreneurs are worse off when they have longer entrepreneurial experience while those without a high school diploma share the same trajectory regardless of experience. Our results suggest that the degree to which entrepreneurial human capital shapes the outside option decreases with education.

Related literature: The existing literature has primarily focused on explicit barriers to entrepreneurship such as financial frictions or market regulation. We contribute to an emerging literature focusing on *implicit* barriers, specifically those posed by irreversible investment and outside options.

The literature on investment irreversibility has argued that when initial investment to businesses cannot be reverted, the entry threshold to entrepreneurship rises (e.g., Abel and Eberly 1996). In direct relation to our work, Tan (2022) argues that entrepreneurial risk is associated with lack of insurance and illiquidity arising from frictional asset resale. He finds that physical investment irreversibility can account for the bulk of entrepreneurial risk. In our context, human capital specificity is analogous to investment irreversibility in that entry into entrepreneurship requires investing in business-specific human capital that is (partially) nontransferable to paid work upon exit. The contribution of our paper is an empirical assessment of the impact of human capital specificity on labor market outcomes.

Our paper also contributes to a recent literature focusing on the option to return to paid work as a determinant of entrepreneurial entry (Dillon and Stanton 2017; Catherine 2022). These papers treat the evolution of the outside option as exogenous, thereby positing that the choice of starting a business does not introduce additional labor income risk. In contrast, our empirical results suggest that human capital specificity is an important ingredient to dynamically evolving outside options and, thus, to entrepreneurial risk.

More generally, our paper contributes to a wide literature on the human capital aspects of entrepreneurship. Bhandari and McGrattan (2021) find that "sweat equity"—human capital associated with running a business—is equivalent to about 120 percent of US GDP. **?** show that many businesses are sold as whole, signifying the importance of intangible assets such as human capital. Linking entrepreneurial activity to human capital, Queiró (2022) finds that more educated entrepreneurs start larger and faster-growing firms while Kozeniauskas (2023) argues that the propensity of high-skill workers to enter entrepreneurship has decreased in recent decades. We contribute to this wide literature by documenting empirical patterns of entrepreneurial experience shaping outside options.

2. Empirical approach

The ideal experiment to measure the evolution of entrepreneurs' outside option would be to compare the outcomes of two *potential* entrepreneurs. These two hypothesized individuals—let us call them Ashley and Blake—are identical in all aspects, including their current employment in paid work. Since they are ex ante identical, they both have the same propensity to start identically productive business ventures. The thought experiment then goes on to posit that Ashley starts their own business while Blake stays in paid work. Blake's subsequent work trajectory serves as a natural measure of Ashley's opportunity cost at any point; furthermore, if Ashley returns to paid work, we can compare their trajectory from the point of return to Blake's to quantify the wage gains or losses of their entrepreneurial stint.

Unfortunately, comparisons in such ideal experiments are not available to us: we cannot observe potential entrepreneurs. Instead, we compare the wage trajectories of *return-entrepreneurs* (former entrepreneurs after returning to paid work) to the trajectories of *never-entrepreneurs* (other workers who never started a business), and we look at the gradient of wage differences by the length of entrepreneurial experience. For example, if a return-entrepreneur of ten years suffers a larger wage loss than a return-entrepreneur of only one year, this suggests that the outside option of returning to paid work worsens over the course of entrepreneurship. We include a rich set of controls to ensure that return-entrepreneurs are similar to never-entrepreneurs, making our comparisons as tight as possible.

We formalize this argument in a regression framework with a number of fixed effects. Assume we obtain panel data on return-entrepreneurs and never-entrepreneurs indexed by i over time t. We observe workers' wages w_{it}, the duration that return-entrepreneurs have spent in paid work dur_{it}, and their demographic and employment characteristics captured by the grouping function $g(i,t)$. Specifically, $g(i,t)$ denotes worker i's gender, age, education, the calendar year at time t, as well as the occupation, sector, and location they work in at time t, and $\varphi_{g(i,t)}$ is the corresponding collection of fixed effects.

Putting all the pieces together, we estimate the following regression specification:

$$\log w_{it} = \sum_{s=1}^{S} \theta_s \mathbb{1}(dur_{it} = s) + \varphi_{g(i,t)} + \varepsilon_{it}. \tag{1}$$

Our parameters of interest are the sequence of θ's which gives us the relative wage trajectory of return-entrepreneurs, compared to never-entrepreneurs, over S years after returning to paid work. The inclusion of the fixed effects $\varphi_{g(i,t)}$ signifies that we make these comparisons within narrowly defined groups. In practice, we will first estimate a model with gender-by-age, education, occupation, and calendar year fixed effects, then iteratively add sector and location fixed effects; the results will reveal composition effects which have important implications for policymakers.

Next we look at the heterogeneity of these trajectories by the length of entrepreneurial experience $exper_i$.[1] We estimate the following regression:

$$\log w_{it} = \beta_0 exper_i + \sum_{s=1}^{S} (\theta_s + \beta_s exper_i) \mathbb{1}(dur_{it} = s) + \varphi_{g(i,t)} + \varepsilon_{it}. \tag{2}$$

Now our parameters of interest are the θ's and the β's: these parameters describe the relative wage trajectory of return-entrepreneurs for a given length of entrepreneurial experience. Note that this specification is linear in experience. Our modeling choice is driven by data limitations: since we only observe workers for 16 years, we can only track return-entrepreneurs with long experience for a few years. The linear specification allows us to make out-of-sample predictions about the relative wage trajectories for the whole length of our sample, even for return-entrepreneurs with 15 years of experience.

3. Data

We implement our empirical approach on the Quadros de Pessoal (QP) dataset which covers the universe of work histories at Portuguese private firms employing paid workers. We can track a person's work history across employers, accompanied by detailed information on wages, occupations, job titles, and demographics like age, gender, and education. We supplement these data with information on the sector and location of firms from the Sistema de Contas Integradas das Empresas (SCIE) dataset. Our final sample covers the time period of 2004–2020.

One key empirical question for our purposes is how we identify entrepreneurs in the data. The QP–SCIE dataset does not allow us to observe them directly, which is commonplace in large-scale administrative data (see also Félix, Karmakar, and Sedláček 2021). We follow Queiró (2022) and define entrepreneurs as top managers of newly established firms. Top managers are identified as (i) directors according to their 4-digit occupation and (ii) employers according to their professional status. Newly established

1. Note that $exper_i$ varies across individuals but not time. Our sample includes return-entrepreneurs only after returning to paid work, so their entrepreneurial experience is fixed.

	Return-entrepreneurs		Never-entrepreneurs	
Num. observations				
Workers	47,611		5,556,832	
Firms	48,194		661,262	
Share of obs. by education				
Less than high school	30.8		59.4	
High school	24.7		24.1	
College	44.5		16.5	
Most frequent occupations [share]				
1	Sales workers	[8.8]	Sales workers	[7.7]
2	Office clerks	[7.7]	STEM occs.	[6.7]
3	Service workers	[7.1]	Administrators	[6.6]
Most frequent sectors [share]				
1	Real estate	[21.2]	Wholesale	[18.6]
2	Wholesale	[20.2]	Real estate	[14.9]
3	Construction	[10.2]	Construction	[9.5]
Statistics (means [25th ǀ 50th ǀ 75th perc.])				
Male (percent)	59.7		53.4	
Age (years)	41.3	[35 ǀ 40 ǀ 47]	40.1	[32 ǀ 39 ǀ 48]
Monthly wage (EUR)	997	[625 ǀ 915 ǀ 1,636]	763	[557 ǀ 728 ǀ 1,074]
Entrep. experience (years)	3.1	[1 ǀ 2 ǀ 4]	–	
Statistics by education (means [25th ǀ 50th ǀ 75th perc.])				
Less than high school				
Male (percent)	73.1		58.1	
Age (years)	41.3	[33 ǀ 41 ǀ 49]	41.2	[33 ǀ 42 ǀ 51]
Monthly wage (EUR)	642	[500 ǀ 623 ǀ 839]	621	[504 ǀ 638 ǀ 825]
Entrep. experience (years)	2.7	[1 ǀ 1 ǀ 4]	–	
High school				
Male (percent)	65.7		50.2	
Age (years)	38.6	[31 ǀ 38 ǀ 45]	36.1	[28 ǀ 35 ǀ 43]
Monthly wage (EUR)	851	[574 ǀ 768 ǀ 1,275]	789	[597 ǀ 766 ǀ 1,104]
Entrep. experience (years)	2.7	[1 ǀ 1 ǀ 4]	–	
College				
Male (percent)	55.9		44.6	
Age (years)	39.2	[32 ǀ 38 ǀ 45]	37.1	[30 ǀ 35 ǀ 43]
Monthly wage (EUR)	1,483	[902 ǀ 1,545 ǀ 2,619]	1,326	[890 ǀ 1,323 ǀ 2,048]
Entrep. experience (years)	3.5	[1 ǀ 3 ǀ 4]	–	

TABLE 1. Summary statistics

Notes: Return-entrepreneurs are paid workers with an observed entrepreneurial history. Never-entrepreneurs are paid workers who are not observed to have started a business in sample. Workers and firms are anonymized. The firm count for return-entrepreneurs shows the number of firms that employ at least one return-entrepreneur. Educational groups are based on 1-digit educational categories (less than high school: did not finish 12th grade; high school: finished 12th grade but did not earn a bachelor's degree; college: earned a bachelor's degree and may have acquired higher levels of education). Occupations and sectors are measured on the 2-digit level. All statistics for return-entrepreneurs are measured after returning to paid work.
Source: QP–SCIE, authors' calculations.

firms are firms whose (anonymized) identifier appears first in the sample after the first observed calendar year.

The resulting sample covers 5.6 million workers across 709 thousand firms (Table 1). 0.8 percent of these workers are return-entrepreneurs, employed at firms making up 6.8 percent of the firm distribution. Return-entrepreneurs are most commonly working

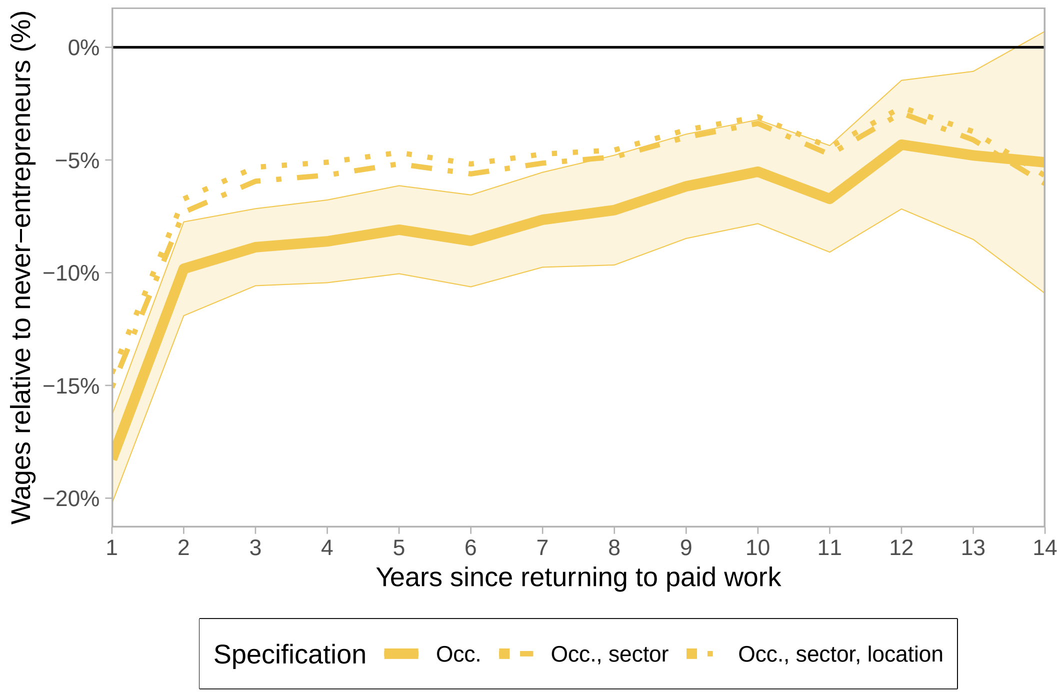

FIGURE 1: Wage trajectories after returning to paid work

Notes: Solid line represents results from specification with occupation FEs. Dash-dotted line represents results from specification with occupation and sector FEs. Dotted line represents results from specification with occupation, sector, and location FEs. Shaded region represents 95 percent confidence bounds around the specification with occupation FEs.
Source: QP–SCIE, authors' calculations.

in sales, office clerk and service occupations while never-entrepreneurs are most likely in sales, STEM, or administrative occupations. Both return-entrepreneurs and never-entrepreneurs are most likely to work in the real estate, wholesale, or construction sector. Return-entrepreneurs are, on average, 6.3 percentage points more likely to be male than never-entrepreneurs, 1.2 years older, 28 percentage points more likely to have graduated from college, and earn 30.5 percent higher wages. Recall, however, that this latter difference is not the relevant comparison: we need to compare return-entrepreneurs with *similar* never-entrepreneurs. Indeed, the wage differences are smaller (4–12 percent) within education groups, so the aggregate wage difference is due to composition effects across education. We make even narrower comparisons in our empirical approach.

4. Results

We now present the results of taking our empirical approach to the data. We first show the estimated relative wage trajectory of return-entrepreneurs compared to never-entrepreneurs. Figure 1 displays the estimated θ coefficients from three alternative versions of Equation 1: one with only occupation fixed effects, one with occupation and sector fixed effects, and one with occupation–sector–location fixed effects.

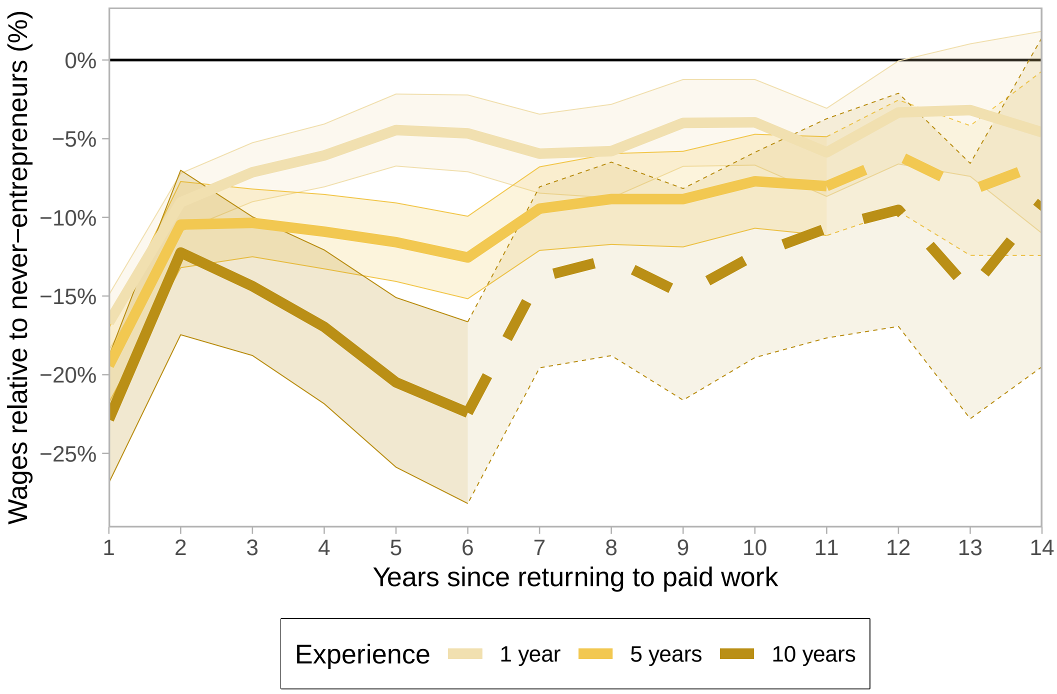

FIGURE 2: Wage trajectories by entrepreneurial experience

Notes: Regression results from specification with occupation FEs. Increasingly darker lines represent wage trajectories after returning to paid work from 1, 5, and 10 years of entrepreneurship, respectively. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

According to our main estimate in Figure 1, return-entrepreneurs—compared to never-entrepreneurs of the same gender, age, education, in the same calendar year, and working in the same occupation—suffer an 18 percent wage loss immediately upon returning to paid work. The wage loss decreases in subsequent years but the trajectory remains flat at a 5 percent loss. The average wage loss is 10.4 percent as shown in Table A.1. Furthermore, the net present value of the wage trajectory is −8.3 percent.[2]

Next, we focus our comparisons on narrower groups. The dash-dotted line in Figure 1 shows the wage trajectory of return-entrepreneurs relative to same-gender-age-education never-entrepreneurs in the same occupation and sector. The immediate wage loss is somewhat smaller but still sizeable at 15 percent, and the long-run wage loss shrinks to 3 percent from 5. The differences between the two estimates are due to composition effects: return-entrepreneurs select into sectors with relatively low wages. Therefore, their wage loss appears to vanish faster when we compare them to the lower baseline of never-entrepreneurs. The same logic applies when we add location to the mix: return-entrepreneurs select into locations with lower than average wages.

Now we know that return-entrepreneurs suffer a wage loss upon returning to paid work, but how does the very act of running a business impact their wages? Figure

2. For this back-of-the-envelope net present value calculation, we set the discount rate to 0.98 and real wage growth to its observed value of 0.05 percent over our sample period 2004–2020.

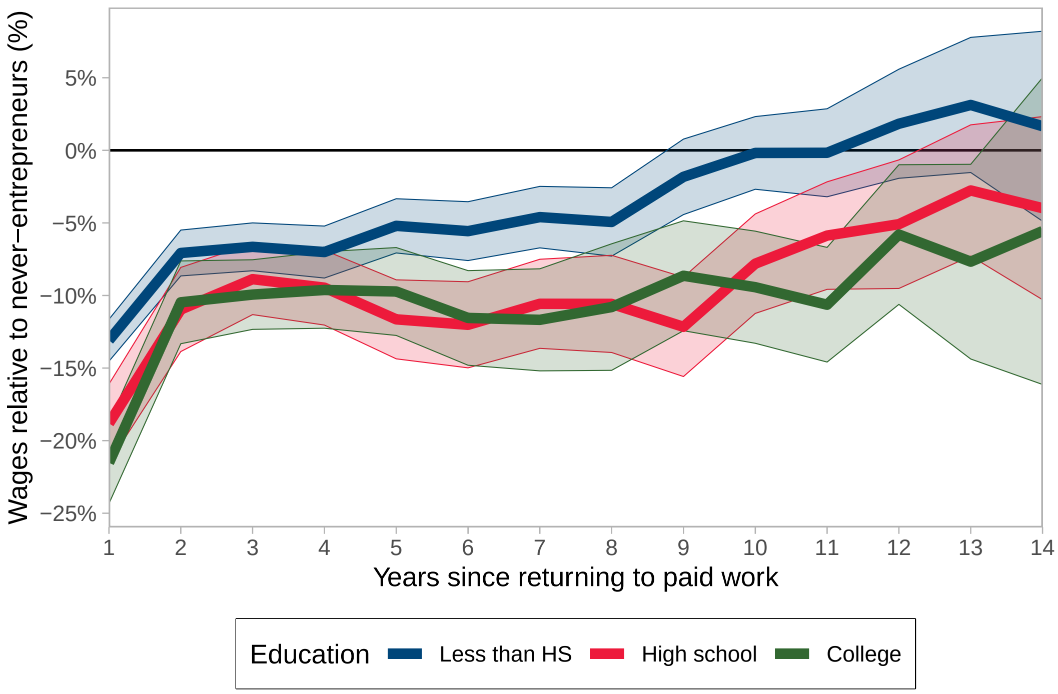

FIGURE 3: Wage trajectories by education

Notes: Regression results from specification with occupation FEs. Blue, green, and yellow lines represent wage trajectories for return-entrepreneurs with less than high school education, high school diploma, and college degree, respectively. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

2 answers this question by displaying the predicted values from three estimates of Equation 2: estimates with 1, 5, and 10 years of entrepreneurial experience.[3] The results are stark: return-entrepreneurs with longer experience fare worse than those with shorter entrepreneurial stints. The wage loss of return-entrepreneurs with 5 years of experience is 3 percentage points larger than those with only 1 year of experience, and the gap grows to 6 percentage points for those with 10 vs. 1 years of experience.[4] These results suggest that time spent in entrepreneurship (i.e., time out of paid work) diminishes the human capital of return-entrepreneurs.

We now turn to discussing how our results differ across return-entrepreneurs by their level of education. We replicate the above analysis separately for three groups of return-entrepreneurs: those without a high school diploma, high school graduates, and college graduates.

3. We only show results from the specification with occupation fixed effects in the main text. The results from the other two specifications are shown in Appendix Figures A.1 and A.2: those estimates are qualitatively similar to the ones presented in the main text.

4. Some of our long-run estimates are out-of-sample predictions since we only observe return-entrepreneurs with 5 (10) years of experience for 10 (6) years. We denote these out-of-sample predictions with dashed lines on our graphs.

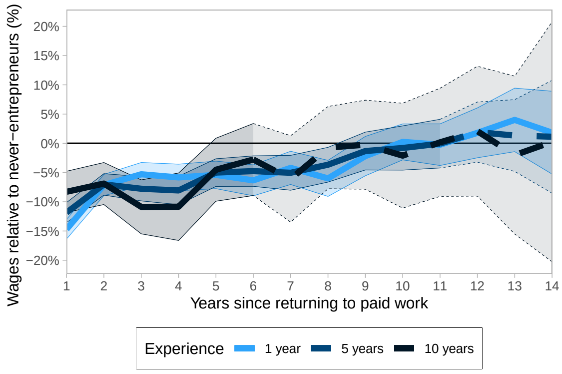

(A) Less than high school

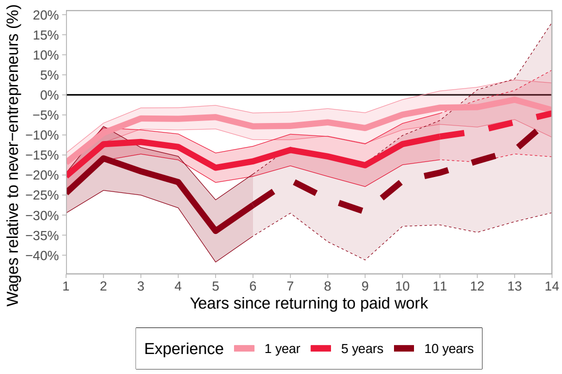

(B) High school

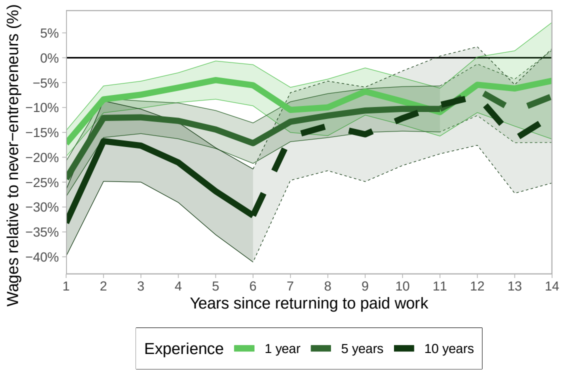

(C) College

FIGURE 4: Wage trajectories by education and entrepreneurial experience

Notes: Regression results from specifications with occupation FEs. Increasingly darker lines represent wage trajectories after returning to paid work from 1, 5, and 10 years of entrepreneurship, respectively, within each panel by education. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

Figure 3 displays the wage trajectories of return-entrepreneurs relative to never-entrepreneurs by education. Return-entrepreneurs with higher levels of education suffer larger and more persistent wage losses. Those without a high school diploma experience a 13 percent wage loss immediately after returning to paid work, and catch up to similar never-entrepreneurs within 10 years. At the same time, those with a high school diploma start from 18 percent lower wages (21 percent if they have a college degree), and do not catch up to their never-entrepreneur counterparts within our sample. These results are in line with entrepreneurship-specific human capital that is imperfectly transferable to paid work. If higher-educated entrepreneurs run businesses with more specific human capital and cannot transfer it to paid work, their wages would presumably be lower upon return. At the same time, if return-entrepreneurs with lower levels of education start businesses with less specific human capital, their wage loss would be less severe. These patters exactly coincide with our results.

As a final step, we take a look at the heterogeneity of our education-specific results by entrepreneurial experience in Figure 4. Return-entrepreneurs with less than high school education do not suffer differential losses by entrepreneurial experience. However, the losses of those with high school or college education are increasing with experience. These patterns further suggest that human capital specificity plays an important role in shaping outcomes after returning to paid work. More educated entrepreneurs who presumably start businesses with more specific human capital suffer larger losses when they ran their business longer; at the same time, entrepreneurial experience does not shift the wage trajectories of less educated return-entrepreneurs.

5. Conclusion

Potential entrepreneurs face many explicit barriers to starting a business. It is costly to set one up, entrepreneurs often need to take on debt to finance their endeavours, and there is uncertainty about how well a business will perform. When businesses close, entrepreneurs often return to paid work. However, accumulating business-specific human capital may affect this outside option, thus posing an implicit barrier to entrepreneurship.

By comparing wage outcomes for entrepreneurs with similar individuals who never started businesses, we show that entrepreneurs initially suffer wage losses when returning to paid work. These wage losses gradually diminish, but it takes years for former entrepreneurs to catch up. Furthermore, both wage losses and recovery to baseline are increasing in entrepreneurial tenure. These effects are also stronger for more educated return-entrepreneurs. Our results suggest that the depreciation of the outside option is an additional barrier to becoming an entrepreneur. Policies that help potential entrepreneurs overcome this barrier, thus insuring against entrepreneurial risk, may improve the allocation of resources, enhance the entrepreneurial environment through business dynamism, and ultimately promote economic growth. Simulating the impact of such policies is an important avenue for future research.

References

Abel, Andrew B and Janice C Eberly (1996). "Optimal Investment with Costly Reversibility." *Review of Economic Studies*, 63(4), 581–593.

Bhandari, Anmol and Ellen McGrattan (2021). "Sweat Equity in U.S. Private Business." *Quarterly Journal of Economics*, 136(2), 727–781.

Buera, Francisco J. and Yongseok Shin (2013). "Financial Frictions and the Persistence of History: A Quantitative Exploration." *Journal of Political Economy*, 121(2), 221–272.

Cagetti, Marco and Mariacristina De Nardi (2006). "Entrepreneurship, Frictions, and Wealth." *Journal of Political Economy*, 114(5), 835–870.

Catherine, Sylvain (2022). "Keeping Options Open: What Motivates Entrepreneurs?" *Journal of Financial Economics*, 144, 1–21.

Dillon, Eleanor and Christopher Stanton (2017). "Self-Employment Dynamics and the Returns to Entrepreneurship." Working Paper.

Félix, Sónia, Sudipto Karmakar, and Petr Sedláček (2021). "Serial Entrepreneurs and the Macroeconomy." CEPR Discussion Paper 16449.

Hopenhayn, Hugo (1992). "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 60(5), 1127–1150.

Kozeniauskas, Nicholas (2023). "What's Driving the Decline in Entrepreneurship?" Working Paper.

Melitz, Marc (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6), 1695–1725.

Queiró, Francisco (2022). "Entrepreneurial Human Capital and Firm Dynamics." *Review of Economic Studies*, 89, 2061–2100.

Tan, Eugene (2022). "Entrepreneurial Investment Dynamics and the Wealth Distribution." Working Paper.

Appendix: Additional results

Dep. var.: log wage	(1)	(2)	(3)	(4)	(5)	(6)
1(return-entrep.)	-0.1036^{***}	-0.0728^{***}	-0.0768^{***}	-0.0471^{***}	-0.0717^{***}	-0.0443^{***}
	(0.0083)	(0.0091)	(0.0058)	(0.0073)	(0.0056)	(0.0071)
Entrep. experience		-0.0101^{***}		-0.0097^{***}		-0.0089^{***}
		(0.0021)		(0.0019)		(0.0019)
Observations	41,793,126	41,793,126	41,793,126	41,793,126	41,793,126	41,793,126
Gender×age FE	✓	✓	✓	✓	✓	✓
Education FE	✓	✓	✓	✓	✓	✓
Calendar year FE	✓	✓	✓	✓	✓	✓
Occupation FE	✓	✓	✓	✓	✓	✓
Sector FE			✓	✓	✓	✓
Location FE					✓	✓
R^2	0.3497	0.3497	0.3912	0.3912	0.3979	0.3979

TABLE A.1. Lifetime wage loss after returning to paid work

Notes: Standard errors, clustered at the firm level, in parentheses. *** p<0.001, ** p<0.01, * p<0.05. Sample includes entrepreneurs after returning to paid work as well as never-entrepreneurs. Outcome variable is log wages. Each year of age has its own fixed effect parsed by gender. Education is measured as 1-digit educational categories. Occupations and sectors are measured on the 2-digit level. Locations are NUTS II statistical regions.
Source: QP–SCIE, authors' calculations.

Table A.1 shows our results on the lifetime wage loss after returning to paid work. The sample includes return-entrepreneurs after their return to paid work as well as never-entrepreneurs. Column (1) shows that the wage loss of return-entrepreneurs is 10.4 percent on average, compared to never-entrepreneurs in the same gender, age, education, and occupation group in the same calendar year. Looking at the dynamics of wage losses, column (2) shows that the baseline wage loss is 7.3 percent and each year of entrepreneurial experience goes along with a further 1 percentage point decrease.

Columns (3)–(4) and (5)–(6) repeat this analysis with the addition of sector and location fixed effects. The results with sector fixed effects in columns (3)–(4) indicate that return-entrepreneurs suffer a 7.7 percent lifetime wage loss on average, with each year of entrepreneurial experience adding 1 percentage point. These estimates imply that return-entrepreneurs return to lower-paying sectors; that is, their wage loss is lower because we compare them to a lower baseline. The same logic applies to the inclusion of location fixed effects in columns (5)–(6). All results are statistically significant and economically sizable.

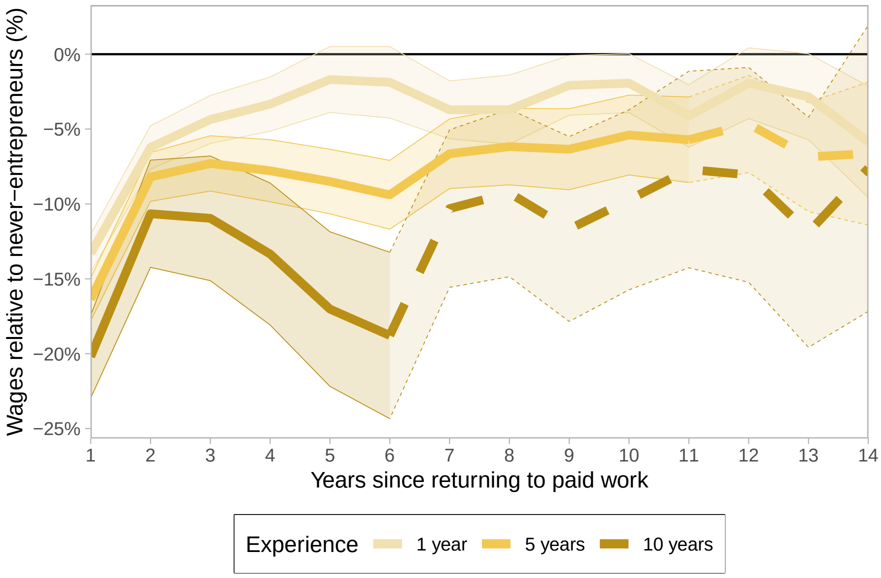

FIGURE A.1: Wage loss by entrepreneurial experience with sector fixed effects

Notes: Regression results from specification with occupation and sector FEs. Increasingly darker lines represent wage trajectories after returning to paid work from 1, 5, and 10 years of entrepreneurship, respectively. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

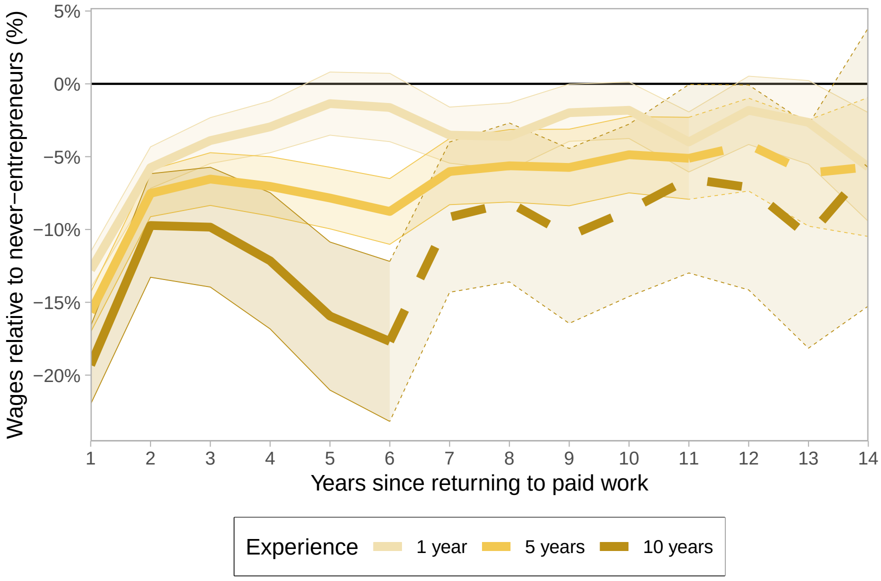

FIGURE A.2: Wage loss by entrepreneurial experience with sector and location FEs

Notes: Regression results from specification with occupation, sector, and location FEs. Increasingly darker lines represent wage trajectories after returning to paid work from 1, 5, and 10 years of entrepreneurship, respectively. Shaded regions represent 95 percent confidence bounds. Dashed line segments represent out-of-sample predictions.
Source: QP–SCIE, authors' calculations.

Figures A.1 and A.2 replicate the results shown in Figure 2 with the inclusion of more fixed effects, resulting in narrower comparisons. The three sets of results are qualitatively similar: longer entrepreneurial experience goes along with more persistent wage losses. When comparing return-entrepreneurs to never-entrepreneurs in narrower groups (same sector and location, not just same occupation), the losses are less persistent. These results, similar to those presented in Figure 1, imply that return-entrepreneurs select into low-paying sectors and locations upon returning to paid work.