N° 547 - 2025

Economie et Statistique

Economics and Statistics

Dossier
Évaluation des politiques publiques
Varia
Discriminations

Thematic Section
Evaluation of Public Policies
Varia
Discrimination

Insee
Mesurer pour comprendre

Economie et Statistique

Economics and Statistics

La revue est en accès libre sur le site **www.insee.fr**. Il est possible de s'abonner aux avis de parution sur le site. La revue peut être achetée sur le site **www.insee.fr**, rubrique « Services / Acheter nos publications ».

The journal is available in open access on the Insee website **www.insee.fr.** Publication alerts can be subscribed online. The printed version of the journal (in French) can be purchased on the Insee website **www.insee.fr**.

Directeur de la publication / Director of Publication:
Fabrice LENGLART
Rédactrice en chef / Editor in Chief:
Dominique GOUX
Rédacteur en chef adjoint / Deputy Editor in Chief:
Riyad ABBAS
Assistante éditoriale / Editorial Assistant: Véronique EGLOFF
Traductions / Translations: FLUENT PLANET
Château Rouge Wood Parc bâtiment C-D 274 ter/3,
Avenue de la Marne, 59700 Marcq-en-Baroeul, France
Maquette PAO / CAP: LUMINESS
1, rue du Docteur-Sauvé, BP3, 53101 Mayenne, France
Impression / Printing: PARAGON TRANSACTION
39 rue des Rivières Saint-Agnan, 58200 Cosne-Cours-sur-Loire, France

Economie et Statistique / Economics and Statistics
Issue 547 – 2025

Preventing Discrimination Through Training Measures: An Assessment

Laetitia Challe*, Sylvain Chareyron, Yannick L'Horty* and Pascale Petit***

Abstract – In France, training of recruiters is often emphasised as an effective means for combating discrimination and in 2017 it was made compulsory for companies with more than 300 employees. In this study, we assess the effect of a measure similar to this compulsory training by comparing the results of correspondence tests performed before and after implementation of the measure in treated companies and control companies. The results show that the level of discrimination was the same between the two groups of companies prior to implementation of the measures and that the same was true five months later. Double- and triple-difference estimates show no significant impact of these measures on the level of discrimination in access to employment.

JEL: J7, C93
Keywords: discrimination, access to employment, correspondence test, awareness raising measures

**Université Gustave Eiffel, ERUDITE (EA 437), TEPP (FR 2042) ; **Université Paris-Est Créteil, ERUDITE (EA 437), TEPP (FR 2042).*
Correspondence*: sylvain.chareyron@u-pec.fr*

This study was funded by the Fonds d'Expérimentation pour la Jeunesse *(Fund for Experimentation for Youth) within the framework of the APDISCRI2 call for proposals issued in 2018 by the French Ministry for Education, Youth and Sport. This study forms part of the DIADEME project (*Dispositifs Innovants Anti Discriminations : Expérimentations Multisites et Évaluation - *Innovative Anti-Discrimination Schemes: Multisite Experimentation and Assessment). We would like to thank the participants in the* Journées Recherche et Action Publique *(REAP 2024) which took place in October 2024 and those at the 19th conference of the* Fédération Théorie et Évaluation des Politiques Publiques *(TEPP) which took place in December 2024. We would also like to thank the two anonymous rapporteurs and Liliane Bonnal, Ahmed Tritah and the editor in chief Dominique Goux for their comments which made it possible to improve the article.*

Received in January 2024, accepted in November 2024. Translated from "Prévenir les discriminations par une action de formation : une évaluation".

Citation: Challe, L., Chareyron, S., L'Horty, Y. & Petit, P. (2025). Preventing Discrimination Through Training Measures: An Assessment. *Economie et Statistique / Economics and Statistics*, 547, 3–16. doi: 10.24187/ecostat.2025.547.2135

Academic research on discrimination has until now focused mainly on its measurement and interpretation, giving priority to a small number of criteria (gender, ethnic origin or place of residence) and areas (mainly the labour market, and more recently the housing market). It more rarely addresses public policies to combat discrimination. In particular, the contemporary experimental assessment methods which have been disseminated in France since the 2000s have been very seldom applied to anti-discrimination measures.

There are a great variety of measures to combat discrimination and promote diversity: a reminder of the law, awareness raising, specific training courses, information given to those who may be victims of discrimination, signing charters and commitments, equality and diversity labels, etc. Training of recruiters is a key aspect of these measures, since it was made compulsory in France for companies with more than 300 employees, by the Equality and Citizenship Act of 29 January 2017. This type of measure has not yet undergone any rigorous impact assessment, however.

In the field of discrimination prevention measures, research has so far mainly concerned the assessment of policies to change the functioning of the labour market by controlling the information available to the recruiter. The introduction of anonymous CVs is one example of this (Behaghel *et al.*, 2015; Krause *et al.*, 2012), while another is the ban on requesting applicants' police records in the United States (Agan & Starr, 2018). This type of policy has so far revealed limited effects or has even proved counter-productive. Some studies have also assessed the effect of incentive policies on discrimination related to place of residence and have shown effects that are positive, although also limited (Chareyron *et al.*, 2022). More coercive measures such as reminders of the law and the threat of legal sanction have also been assessed, mainly in the housing market (Chareyron *et al.*, 2023; Fang *et al.*, 2019; Murchie *et al.*, 2021).

Given these mixed results, it seems a good idea to examine the effect of measures at company level, particularly in the recruitment process. For example, Berson *et al.* (2020) showed that access to a centralised human resources department in the recruitment process tended to reduce discrimination against applicants of North African origin.

We propose to assess the effects of a recruiter training measure on discrimination based on ethnic origin in access to employment. In its content and intensity, the measure is similar to the compulsory training provided for by the French Labour Code (Art. L1131-2). It was implemented by the Regional Association of Local Youth Employment Centres (*Association régionale des missions locales*) of the Provence-Alpes-Côte d'Azur (PACA) region, which provides access to a one-day training session for a number of volunteer partner companies, whether or not they are concerned by the legal obligation, in order to raise their awareness of the issue of discrimination in recruitment. In this study, we also address a second experiment with similar content that was rolled out in a different region, by the Local Youth Employment Centre (*Mission locale*) of the town of Chambéry. These training courses have legal and practical content, stressing the legal and judicial framework of the prohibition of discrimination and the procedures for organising recruitment in conformance with the principle of equality.

To carry out the assessment, we measure discrimination risk by the correspondence test experimental method. We make a comparison between companies in which the training measure has effectively been implemented and others where it has not been rolled out. The companies in the control group included in the experiment are selected on the basis of their observable characteristics in terms of sector of activity and location. An initial test is performed in the month prior to the training course and a second test within the four months post-training. We test the people who took part in the training directly. A difference-in-differences estimation then enables us to eliminate the time-invariant unobserved differences between the two groups and the variations over time common to both groups.

The results show a substantial and comparable level of discrimination in the companies of the treated and control groups in the period prior to the treatment. In the control group, the applicants whose surname and first name suggest a North African origin have about a 12.5 percentage-point lower probability of receiving a positive response to their application than an applicant whose surname and first name suggest a French origin. This represents a relative difference of nearly 50%. The difference is the same five months later and apparently was not affected by implementation of the measures. The double- and triple-difference results show no significant effect of the training measures on the level of discrimination. This suggests that the public policy consisting of imposing a compulsory training measure of low intensity with mainly legal content is not an adequate response to the challenge of ethno-racial discrimination in the labour market.

Section 1 describes the support measures implemented for companies, while Section 2 describes the experiment protocol for data collection. Section 3 describes the empirical strategy, while Section 4 presents the results of the estimates, before the conclusion.

1. A Training Measure to Combat Discrimination

1.1. Content of the Measure

One of the two training measures assessed in this article was implemented by the Regional Association of Local Youth Employment Centres (*Association régionale des missions locales*) of the PACA region. It consists in proposing to a number of volunteer companies to receive training which involves legally defining non-discrimination and its fields of application based on the law and established legal precedents, providing an objective view of a recruitment process that is formally equal, while outlining the principle of company neutrality. The training lasts one day and is intended for recruiters and the people involved in the recruitment circuits, managers or human resource representatives. The aim is to propose a practical recruitment method to prevent discrimination in the recruitment process from a constructive and non-stigmatising perspective focused on the HR policy of the companies. In-person learning was used throughout the experiment. The trainees are managers or representatives of human resources.

The assessment also includes a second measure of a similar duration and content implemented by the Chambéry Local Youth Employment Centre, which proposes training and support to persons in charge of human resources in the partner companies, with a view to rethinking the organisation of recruitment based solely on the applicants' abilities, producing job descriptions based on objective criteria and the tasks and competencies required, and avoiding the use of potentially discriminatory characteristics of the applicants, such as the first name, surname or place of residence. The Chambéry scheme consists in supporting the recruiting companies and young people prior to contact with the companies. The support consists of co-designing new recruitment processes.

The measures of the Local Youth Employment Centre include:

- the in-situ design of a job description based on objective criteria and the tasks and competencies required;

- a commitment by the company to receive one or more of the prospective applicants without going through the conventional stages (CV, cover letter);

- no information on potentially discriminatory characteristics of the applicants (first name, surname, place of residence) until the principle of an interview appointment has been agreed on by the company;

- the development of forms of innovative collective recruitment measures (atypical job dating, "improbable" encounters, employment afterwork sessions, recruitments organised in the local centre).

The training measures that we assess are similar in nature and intensity to the compulsory training the principle of which is defined in Act No. 2017-86 of 27 January 2017 on equality and citizenship and as incorporated into the French Labour Code in Article L. 1131-2.[1] This stipulates that "in every company employing at least three hundred employees and in every company specialised in recruitment, the employees in charge of recruitment tasks shall receive training in non-discrimination in recruitment at least once every five years". To comply with this legal obligation, vocational training actors generally propose one-day in-person or e-learning sessions. The content of the training may vary from one service provider to another, but generally includes an introduction to the legal and judicial framework of discrimination and practical advice for organising a recruitment operation complying with the principle of equality. The courses are mostly intended for managers and HR personnel.

1.2. Expected Effects

Traditionally, two economic explanations have been given for discrimination. The first is referred to as taste-based discrimination and was formally defined by Becker (1957). It is the existence of a preference for employment of people from the majority group in the employer's utility function. In this case, the employer may be prepared to recruit a less productive employee and hence accept a reduction in their profit to avoid employing a person from another demographic group. The second explanation relates to the concept of statistical discrimination formulated by Arrow (1974). In this case, faced with imperfect information on the applicants,

1. *A recent DARES research document entitled «* Quelles sont les caractéristiques de l'offre de formation à la non-discrimination à l'embauche ? *» (March 2024) presents an overview of the offer of training in non-discrimination in recruitment since the enactment of the Equality and Citizenship Act of 2017 (Benedetto-Meyer, 2024).*

employers presume that the unobserved characteristics of applicants from another demographic group are less advantageous on average than those of the candidate from the majority group. To maximise their profit, out of applicants having the same observed characteristics, the recruiter will favour the one that belongs to the majority group. More recently, there has been another trend suggesting the explanation that recruiters can have unconscious prejudices, even if they possibly endeavour to correct them consciously. According to this explanation, discrimination may be greater in the case of heavy cognitive loads or inattention to the task (Bertrand *et al.*, 2005).

From a theoretical standpoint, a training and awareness raising measure regarding the issue of discrimination in recruitment seems an appropriate solution for combating discrimination. Referring back to the above classification, the aim is mainly to combat implicit discrimination due to the pitfalls of stereotypes. Compulsory training offers an opportunity for professionalisation of the HR function in which diversity is a performance issue: recruiting the person with the required competencies. The "best recruitment practice recommendations" part of the training courses addresses the concept of cognitive biases which can be minimised and thereby reduce implicit discrimination. Given that the training measures consist mainly in a reminder of the legal and judicial framework of discrimination and advice on organising a recruitment operation that complies with the principle of equality, it is unlikely that they would affect direct and rational discrimination. Such discrimination is based on the recruiters' biases and their imperfect knowledge of the applicants' productivity levels. For that, the training measures would have to change the biases and/or reduce information asymmetry during recruitment, which seems unlikely.

2. The Experiment Protocol

2.1. Description of the Protocol

Our approach involves comparing exposure to discrimination in the companies of the treated group before they benefit from the training measure, with a sample of comparable companies which do not benefit from the measure, and then reiterating the comparison after implementation of the training measures. To make these comparisons, we use experimental data collected by the repeated correspondence test method. We accordingly perform two series of tests in each treated company and each control company, the first in the month prior to the measure and the second in the following four months, as described in the diagram below (Figure I).

We created the identities of two fictitious job applicants of the same gender, indicating a French-sounding first name and surname for one and a North African-sounding first name and surname for the other. The first names and surnames were chosen from the most common ones in France based on the official registry files managed by INSEE.

The experimenter sent the assessor the occupational contact details of the employees who would soon be benefiting from the support measure as they were received. For each treated company one or more control companies were then identified (same sector of activity and same employment zone).[2] We then identified an occupation that is commonly recruited in

2. The control companies were tested in the same time frame as the companies of the test group. We used several websites such as societe.com or infogreffe.fr which list the companies still in activity in a given sector and a given region. In the whole database, we identified 93 different sectors of activity (5-figure NAF activity code), ranging from construction to industry, and including market and non-market services. With regard to geography, insofar as possible we compared the control companies with the companies of the treated group (either in the same municipality, or in a nearby geographic area still remaining within the same department). Information on the size of the companies is not always present on these sites.

Figure I – **Experiment protocol**

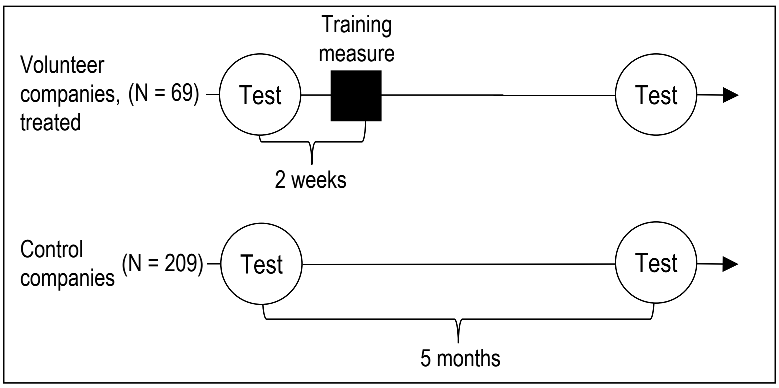

the sector of activity of the treated company and which appears among the in-demand occupations (source: INSEE).[3] For this occupation, we adopted the modal gender of the employees based on job survey data. The occupation was changed between the two series of correspondence tests to minimise the risk of detection of the experiment. The number of control companies for each treated company varies (between 1 and 9) depending on the sector of activity of the treated company, and depending on whether it was more or less easy to find control subjects meeting the sector and location criteria. This explains the variation in the number of control companies between tests.

Each of the two fictitious applicants sent the treated company and the corresponding control companies a message requesting information on employment opportunities in the selected occupation. This is therefore a test similar to a spontaneous application (Deuchert & Kauer, 2017). The spontaneous application test is one of the two types of tests most commonly used in the economic literature on discrimination, together with the test in response to job offers (Riach & Rich, 2010).

According to the OFER survey (*Offre d'emploi et recrutement*, Job Offer and Recruitment) survey conducted by the French Ministry of Labour, 21% of recruitments result from a spontaneous application, and in 68% of recruitments, spontaneous applications were examined during the recruitment procedure. Spontaneous applications are therefore a recruitment channel that should not be overlooked. Moreover, certain companies now have a spontaneous application form on their career website, suggesting to candidates to apply in this way.[4] The development of CV libraries also contributes to the role assigned to spontaneous applications as a potential recruitment channel. In this case, however, no CV or cover letter was sent. The test can therefore be considered as a test by information request. Below are examples of messages sent by two fictitious applicants to a given company for a secretarial job.

> *Hello,*
>
> *I would like to apply for a job in your company. I am a secretary. Could you provide me with information regarding the opportunities and the person to contact? Thanking you in advance for the information you are able to provide.*
>
> *Kind regards,*
> *Jamila BELHADJ*

> *Hello, I am looking for a secretarial job. Could you tell me whether there might be opportunities in your company and, if so, who I can send my application to? Thanking you in advance.*
>
> *Kind regards,*
> *Aurélie Legrand*

The order in which the two requests for information were sent to a given company was determined randomly. Two precautions were taken to limit the risk of detection. First, during series 1 and series 2, several days went by between sending the two messages to the same company. Next, the identities of the fictitious applicants were changed regularly. In particular, they were different in series 1 and series 2. The recruiter's response was considered negative if they did not reply to the message requesting information or if they stated explicitly that there was no job opportunity. The response was considered positive if the employer requested additional information (CV) or put the applicant in contact with the HR Department, etc. The responses that the companies gave to the requests from the two fictitious applicants were then compared.

2.2. Sample Size, Attrition and Data Balancing

The data were collected between June 2019 (first tests of series 1) and May 2023 (final tests of series 2). Collection was suspended during the first lockdown (between March and May 2020). 213 companies were included in the treated group and 629 in the control group. The assessment covered the companies that were able to be tested twice, before and after the training measure. If a company in the treated group was unable to be tested twice, all the companies of the control group associated with it were deleted. After removing certain companies in the PACA region, which were due to take part in the training but which did not do so in the end, and the associated companies of the control group, the final sample comprises 202 companies in the treated group (133 for the PACA region and 69 for Chambéry) and 606 companies in the control group (397 for the PACA region and 209 for Chambéry). The treated companies and control companies are identical in series 1 and series 2.

3. We count in all 28 tested occupations which are cross-cutting to most of the sectors, such as administrative employee, accounting employee, non-specialist salesperson, production worker, and unskilled technician.
4. Over our field of study, about 30% of the companies of the control group had this type of form. This proportion is fairly similar for very small companies (34%) and large companies (30%), but lower for companies of unknown size (21%). In this test, the applications were made via this type of form when it existed, or by email when it did not.

The email addresses are usually those of individuals in the companies of the treated group. In the case of Chambéry, for example, we have 42 non-generic email addresses for the 69 test companies, versus 8 non-generic addresses for the 209 control subjects. A non-generic email address assures us that we are testing the same individual during the two series of correspondence tests. This is therefore the case for 60% of the companies of the treated group (versus 4% for the companies of the control group). This difference of proportion is due to the fact that the control companies can mostly be reached by filling in a contact form without knowing beforehand the identity of the person who will actually reply to the request for contact.

2.3. Estimation Strategy

As we selected control companies that were similar in terms of their sector and location for each company in the treated group, the two groups of companies have fairly similar characteristics. However, to avoid companies in the control group receiving the training measures stipulated by law for companies with more than 300 employees during the experimentation period, we were unable to select control group companies of similar size to those of the treated group. In the absence of random selection of the companies of the control group, there could also be unobserved differences in characteristics between the two groups. If those characteristics affect the result variable (i.e. the fact that an applicant receives a positive response to their application), estimating the effect of the treatment by comparing the differences in positive response rates between the applicant presumed of North African origin and the one presumed of French origin in the two groups while controlling for the observable characteristics would lead to biased results. Another possibility would be to observe the variation in differences in the positive response rates between applicants of North African origin and French origin in the treated group before and after treatment. However, this estimate could be biased by any change that is unrelated to the treatment (in particular the economic situation), which could have different effects on the positive response rate for the applicant of North African origin and for the applicant of French origin.

We therefore adopted a difference-in-differences (or triple differences) estimation strategy. By comparing variations in differences in the positive response rates between applicants of North African origin and applicants of French origin in the treated group and the control group between the period prior to treatment and the period following it, we eliminate the time-invariant unobserved differences between the two groups and the variations over time common to both groups. This triple difference is estimated using the following model:

$$\begin{aligned} REP_{iet} = \alpha + \beta\, NAfr_i + \varphi\, Post_t + \gamma\, T_e \times NAfr_i + \tau\, T_e \\ \times Post_t + \omega\, NAfr_i \times Post_t + \delta\, T_e \times NAfr_i \\ \times Post_t + \pi\, X_{iet} + \mu_t + \phi_e + \varepsilon_{iet} \end{aligned} \quad (1)$$

where REP_{iet} is a dichotomous variable indicating whether company e responds positively to applicant i at date t. $NAfr_i$ is a dichotomous variable indicating whether the applicant is of North African origin. T_e is a dichotomous variable taking value 1 if the company belongs to the treated group and 0 otherwise. $Post_t$ is a dichotomous variable taking value 1 in the post-treatment period and 0 in the period prior to treatment.[5] X_{iet} corresponds to the order in which applications i are sent to company e at date t. μ_t are the fixed effects related to the date of sending the application (month × year and day) and ϕ_e are the company fixed effects. The associated coefficient δ associated with variable $T \times NAfr \times Post$ captures the effect of the training measures.

This strategy is based on the assumption that the variations in differences in the positive response rates according to ethnic origin (i.e. discrimination against the applicant of North African origin) would have been similar in the companies of the treated group and the control group in the absence of treatment (common trend assumption). We cannot observe the trends prior to treatment in order to confirm this assumption. Moreover, inclusion in the treatment is not random and mainly concerns companies that are part of the local youth employment centre's network and are volunteers to take part in the measure, so they are possibly already more aware of discrimination and the need for a change of practices.

However, several factors support the idea that this assumption is complied with. Firstly, while the companies do volunteer to take part in this support, the recruiters who have undergone training as part of this support were generally not the decision-makers on this. From the viewpoint of a recruiter in a treated company, the training is imposed by the company and not a matter on which they are free to make their own decision. Secondly, compliance with this assumption is confirmed by the selection in the

5. Since the treatment was administered as the occasion arose, it does not coincide with a particular date.

control group of companies similar to those in the treated group from the sector and location viewpoints. Moreover, as may be noted below, the companies in the two groups have similar discriminatory behaviours in series 1, even in terms of level. They are therefore unlikely to have seen variations in the level of discrimination that were not related to the support measures in the five months following the first test.

Despite this, a shock that might affect the treatment of applicants of French origin differently to applicants of North African origin in the treated companies and in the control companies could create a bias in the estimate. Given that the second series of tests occurred after the COVID pandemic, it is possible that a shock could have had different effects on the positive response rates of the treated companies and control companies, e.g. because their average size is not the same. On the other hand, it is less likely that this same shock could have affected differently the positive response rate received by the applicant of French origin compared with that obtained by the applicant of North African origin. The latter type of shock would, in theory, be due rather to a political event or to (domestic or external) conflicts liable to alter the recruiters' preferences. This type of shock apparently did not occur between June 2019 and May 2023.

One important point to watch is the possibility that companies may have benefited from other training measures in addition to that provided for in the experiment, especially because of the requirement instituted by the Act of the 2017 that companies with more than 300 employees must propose training on non-discrimination in recruitment for the employees in charge of recruitment tasks.

We cannot know whether the companies in the treated and control groups had implemented the compulsory training at the time of the experiments. However, since the control group largely consists of companies with less than 300 employees, it is unlikely that they received this training during this period.

Moreover, the recruiters of the treated companies may have possibly already taken part in a training course since the introduction of compulsory training in 2017, but it is unlikely that a training course of this type was implemented in addition to the training assessed during the study period, given the frequency stipulated by law. However, we are unable to determine whether there is an equivalence or a complementarity between the compulsory training and the training under the experiments.

Given that the number of companies selected in the control group may vary for each treated company, in the estimates we weight the control group observations by the inverse of the number of control subjects selected for each treated company. Furthermore, the standard errors are clustered at the company level.

3. Results

3.1. Comparison of the Treated and Control Groups Prior io Treatment

In Table 1 we compare the characteristics of the companies in the treated group and the control group for the PACA region and Chambéry territories separately. The differences are generally slight, except on the size of the companies. More specifically, the companies in the treated group are large significantly more often than those in the control group. This under-representation of large companies in the control group is due to the fact that we gave priority to selecting companies with less than 300 employees in the control group to avoid those subject to the training obligation by the terms of the Act of 2017.[6]

Generally speaking, less than a quarter of the companies are in the public sector in Chambéry and about 13% in the PACA region. A large number of sectors of activity are represented. Of the sectors of activity, it may be noted that jobs in general government, the building sector and retailing account for a fairly large proportion of the offers. In the PACA region, temporary work agencies and domestic help companies are also very largely represented. For the measure implemented in Chambéry, 90% of the companies are located in the Savoie department and for the measure implemented in the PACA region, about 30% of the companies are located in Marseille. Apart from the size of the companies, no significant difference of characteristics can be observed between the companies in the treated group and the control group.

The gross rates of response before and after treatment are presented separately in Table 2 for the group of treated companies and the group of control companies. The positive response rates are fairly high by comparison with those of studies using spontaneous application tests (Chareyron *et al.*, 2024) and reach comparable,

6. Less than 5% of the companies in the control group have a size exceeding 300 employees. When it was hard to find a company in the same sector and the same location, a company with more than 300 employees was selected. It is important to note that we do not know the size of all the companies. However, even among those for which we know the size, less than 10% have more than 300 employees. Moreover, it is likely that the companies whose size was unable to be determined are mostly companies with less than 300 employees.

Table 1 – **Comparison of the characteristics of companies in the treated group and the control group**

	Chambéry			PACA		
	Control %	Treated %	Difference	Control %	Treated %	Difference
Public-sector companies	21.5	15.9	−5.6	13.1	12.0	−1.1
Size:						
Large	4.8	15.9	11.1***	4.0	11.3	7.3***
Small or medium	25.8	14.5	−11.3**	24.9	33.1	8.2***
Very small	27.8	17.4	−10.4**	29.0	16.5	−12.5***
Unknown	41.6	52.2	10.6**	42.1	39.1	−3
Sector of activity:						
Temporary work agency	0.0	0.0	0	17.1	12.0	−5.1
Domestic help	1.4	1.4	0	7.6	7.5	−0.1
Public administration	9.6	7.2	−2.4	9.1	7.5	−1.6
Business support	1.9	1.4	−0.5	1.5	1.5	0
Childcare	6.2	4.3	−1.9	3.8	5.3	1.5
Medical accommodation	5.7	4.3	−1.4	5.5	7.5	2
Building	11.5	11.6	0.1	5.0	5.3	0.3
Trades and crafts	8.1	10.1	2	1.0	1.5	0.5
Restaurants and catering	3.8	2.9	−0.9	1.5	0.8	−0.7
Retailing	14.8	18.8	4	10.1	8.3	−1.8
Tourist accommodation	0.0	0.0	0	3.8	3.0	−0.8
Social welfare	3.3	2.9	−0.4	1.8	3.0	1.2
Training	5.3	4.3	−1	4.5	3.0	−1.5
Other	28.2	30.4	2.2	27.7	33.8	6.1*
Location:						
Savoie	90.9	92.8	1.9	0.0	0.0	0
Marseille	0.5	1.4	0.9	33.5	34.6	1.1
Observations	418	138		794	266	

Note: *** p<0.01; ** p<0.05; * p<0.1. The p-values are obtained from equality of proportion tests. The information on the company's size is collected using several website such as societe.com and infogreffe.fr. This information is not always present on these sites.
Source: Diademe testing, TEPP-CNRS.

or even higher levels than those observed in certain studies using applications in response to an offer (Edo & Jacquemet, 2013). However, they are slightly lower in the control group than in the treated group. Before and after treatment, the difference in the positive response rates between applicants of North African origin and those of French origin is significant at the 5% threshold in both the treated and control groups.[7] In both groups, discrimination does not vary significantly between the two periods. The before-after variation in discrimination between the two groups, presented in the last column of the table, is not significant either.

Tables A1 and A2 in Appendix show these same results separately for Chambéry and PACA. The results are generally similar in both regions. However, we can see a significant reduction in discrimination in the control group, even though this does not result in a significant difference in the variation between the two groups.

Figure II presents the estimated levels of discrimination for the companies of the treated group and the control group in the PACA region, in Chambéry and for both the regions in the period prior to the start of treatment. The differences in the positive response rates between the applicant of North African origin and the applicant of French origin are significantly different from 0 at the 5% threshold in each of the regions and

7. The positive response rates generally show a declining trend between the two periods. This decline is significant at the 10% threshold for the applicants of North African origin in the treated group and at the 5% threshold for the applicants of French origin in the control group. However, these variations do not significantly change the differences in the positive response rates between applicants.

Table 2 – **Positive response rates before and after treatment**

	Before treatment	After treatment	Difference	Difference-in-differences
Treated group				
French origin	45.54	41.09	−4.45	
North African origin	34.65	26.73	−7.92*	
Difference	−10.89**	−14.36***	−3.47	
Control group				−9.07
French origin	33.66	27.72	−5.94**	
North African origin	19.97	19.64	−0.33	
Difference	−13.69***	−8.08***	5.61	

Note: *** p<0.01; ** p<0.05; * p<0.1. The p-values are obtained from equality of proportion tests. The last column of the table shows the difference in variation of discrimination between the treated group and the control group and between the two periods (−3.47−5.61).
Source: Diademe testing, TEPP-CNRS.

each of the groups except in the treated group for the PACA region where the difference is significant at the 10% threshold. However, the levels of discrimination are not significantly different between the companies of the treated group and the control group.

This lack of any significant difference in the level of discrimination between the two groups suggests that, although they are of different sizes, the companies in the treated group and the control group share similar discriminatory characteristics and behaviours. A comparison of the differences in the positive response rates between the applicant of North African origin and the applicant of French origin after the treatment occurred could therefore give unbiased results. However, it seems wiser to eliminate any unobserved differences by comparing the variations in differences in the positive response rates.

3.2. Difference-in-Differences

Table 2 also gives an initial general view of the variations in positive response rates obtained

Figure II – **Comparison of discrimination levels between the treated group and the control group prior to treatment**

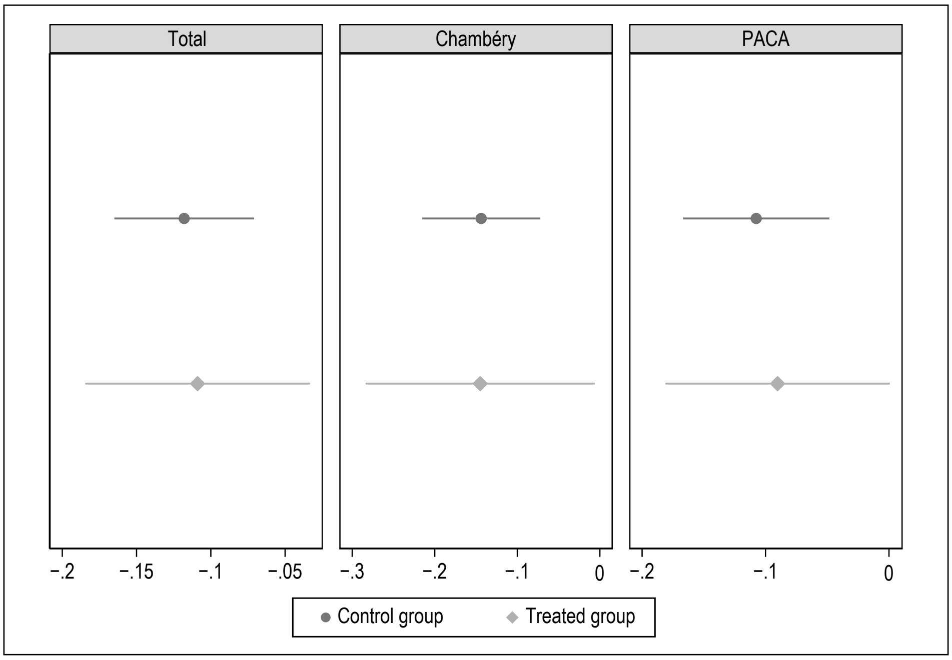

Note: The 95% confidence intervals are presented. The standard errors used to calculate the confidence intervals are clustered at the company level. The control group observations are weighted by the inverse of the number of control subjects selected for each company treated. The estimated $\widehat{\beta}$ coefficients are obtained based on ordinary least squares estimates of the model $REP_{iet} = \alpha + \beta\ NAfr_i + \varepsilon_{iet}$ during the period prior to treatment.
Source: Diademe testing, TEPP-CNRS.

by the two applicants between the period prior to treatment and the post-treatment period. The positive response rate for the applicant of French origin decreases significantly between the two series in the control group but not in the treated group. The positive response rate for the applicant of North African origin decreases, although not very significantly, in the treated group but not in the control group. These facts suggest no substantial reduction in discrimination in the treated group by comparison with the control group.

Table 3 presents the ordinary least squares estimates of equation (1). Columns (1) and (2) present the estimates performed on the whole sample, columns (3) and (4) on the observations relating to the Chambéry region and columns (5) and (6) on the observations relating to the PACA region. In columns (2), (4) and (6), the order in which the requests for information are sent is included in the checks without any major change in the results.[8]

The results show the existence of significant discrimination in the control group in the period prior to the treatment. The applicant whose surname and first name evoke a North African origin has a 12.5 percentage-point lower probability of obtaining a positive response than the applicant presumed of French origin, in the first test in the control group. Since the positive response rate is about 27% in the control group in the period prior to the treatment, in relative terms this corresponds to a penalty of slightly less than 50% to the disadvantage of the applicant of North African origin. This penalty is of the same order of magnitude, although at the high end of the range, as the results obtained by prior studies examining discrimination related to ethnic origin in access to employment. In France, recent studies show that the penalty adversely affecting applicants of North African origin is around 40% (Chareyron *et al.*, 2022). In the United States, Bertrand & Mullainathan (2004) found a 33% difference in the response rate between a white applicant and an Afro-American candidate. In Belgium, Baert *et al.* (2015) obtained 31% fewer responses for the Turkish applicant than for the Flemish applicant. This suggests that the use of requests for information in a correspondence test gives similar results to those obtained with responses to job offers.

However, no significant variation in the level of discrimination can be seen post-treatment or between the control group and the treated group, at the 5% threshold. In particular, the estimated coefficient associated with the variable $T \times NAfr \times Post$, which captures the effect of the training measure, is not significant at the 5% threshold.

However, the lack of significance of the estimated coefficient could be due to a lack of statistical power rather than a real lack of effect of the treatment. Consequently, we now endeavour to determine whether these results can be interpreted as the lack of effect of the training measures. To do so, we calculate the 95% confidence intervals. The results are presented in the second part of Table 3. While it seems unlikely that the effect of the training could be negative, i.e. that it could increase the level of discrimination, considering the upper bound of the confidence interval, we can reasonably rule out an effect of more than 13.3 percentage points. With a 13.7 percentage-point difference in the positive response rates in the control group, this corresponds to a maximum reduction in discrimination of around 100% in relative terms. The estimated effect of the measures is therefore imprecise and it is not possible to rule out a non-negligible effect of training on discrimination.

As mentioned earlier, to take part in the controlled experiment we selected, for each company of the treated group, control companies that are similar in terms of their observed characteristics. The two groups therefore share similar discriminatory characteristics and behaviours prior to the occurrence of treatment. One alternative to the triple differences estimation strategy could therefore be to compare the differences in the positive response rates between the applicant of North African origin and the applicant of French origin after the treatment occurred. Similarly, insofar as it seems unlikely that a shock may have affected the positive response rates of the applicants of North African and French origin differently between the two periods, another possibility could be to compare the change in differences in the response rates for the two applicants between the period prior to the training measure and the period following it. While it does have advantages in terms of identification, triple differences estimation has the disadvantage of having less statistical power.

The results obtained in column (2) of Table 1 are therefore compared with those obtained by double differences on a cross section (post-treatment treated/control comparison) and double differences in time (treated group before/after comparison). Figure III shows the

8. *Deleting from the sample the few companies of the control group with more than 300 employees does not change the results (see Table A3 in Appendix).*

Table 3 – **Effect of the training measures**

Variables	(1)	(2)	Chambéry (3)	Chambéry (4)	PACA (5)	PACA (6)
NAfr	−0.126***	−0.125***	−0.167**	−0.166**	−0.077	−0.089*
	(0.039)	(0.039)	(0.073)	(0.073)	(0.053)	(0.053)
T × NAfr	0.019	0.022	0.047	0.045	−0.031	−0.008
	(0.058)	(0.059)	(0.114)	(0.113)	(0.084)	(0.083)
Post	0.173	−0.064	−0.265	−0.290	0.098	−0.964***
	(0.183)	(0.227)	(0.299)	(0.361)	(0.262)	(0.340)
NAfr × Post	0.042	0.043	0.064	0.064	−0.025	0.003
	(0.053)	(0.053)	(0.087)	(0.086)	(0.071)	(0.072)
T × Post	0.006	−0.008	0.120	0.118	−0.056	−0.106
	(0.062)	(0.061)	(0.113)	(0.112)	(0.085)	(0.086)
T × NAfr × Post	−0.033	−0.040	−0.078	−0.078	0.065	0.032
	(0.084)	(0.085)	(0.146)	(0.146)	(0.110)	(0.109)
Lower CI	−0.198	−0.207	−0.364	−0.364	−0.151	−0.182
Upper CI	0.132	0.127	0.208	0.208	0.281	0.246
Max. relative effect	0.96	0.92	1.50	1.50	2.06	1.81
Date fixed effects	X	X	X	X	X	X
Company fixed effects	X	X	X	X	X	X
Order of sending		X		X		X
Observations	3,232	3,232	1,112	1,112	2,120	2,120
R^2	0.162	0.166	0.293	0.293	0.177	0.195

Note: *** p<0.01; ** p<0.05; * p<0.1. The coefficients estimated by the ordinary least squares method are presented in the table. The standard errors clustered at the company level are shown in brackets. The control group observations are weighted by the inverse of the number of control subjects selected for each company treated. The maximum relative effect is obtained by dividing the upper bound of the confidence interval by the level of discrimination in the control group in the period prior to treatment.
Source: Diademe testing, TEPP-CNRS.

Figure III – **Comparison of the estimated coefficients obtained based on double- and triple-difference estimates**

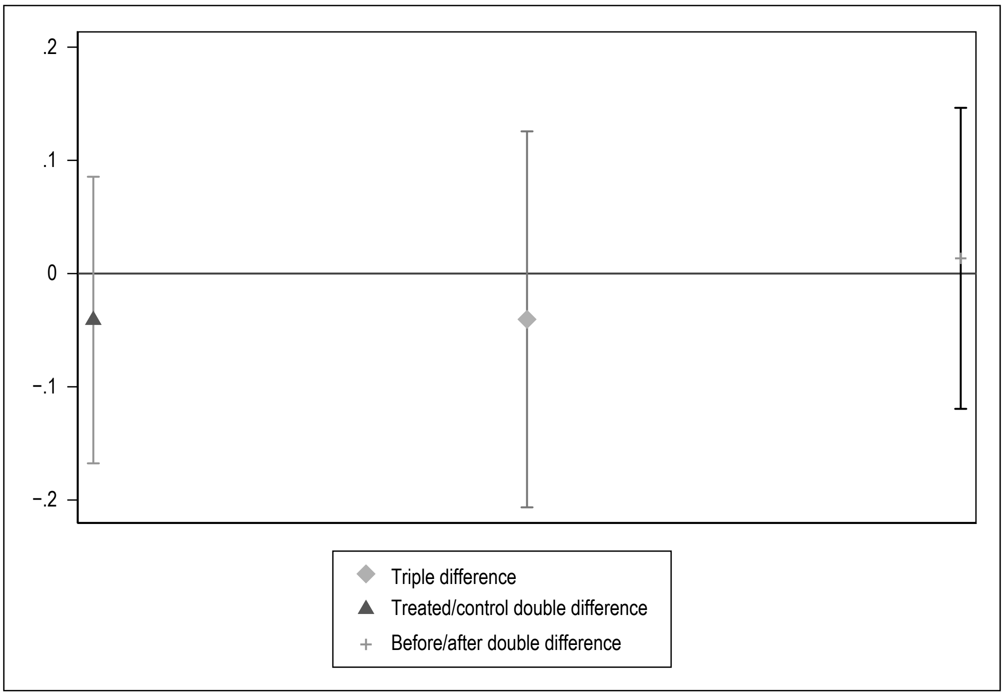

Note: Triple difference corresponds to the estimated coefficient associated with the variable $T \times NAfr \times Post$ of equation (1). Treated/control comparison corresponds to the estimated coefficient associated with the variable $T \times NAfr$ of equation (1), estimated on the sample of companies tested post-treatment. Before/after comparison corresponds to the estimated coefficient associated with the variable $NAfr \times Post$ of equation (1), estimated on the sample of companies belonging to the treated group.

estimated effects of treatment based on these three different strategies. In the three cases, the estimated effect is not significant at the 5% threshold and the estimated coefficients are generally very close to 0. However, although the estimation by comparison of the level of the treated group with the control group is slightly more precise than the others, the upper bounds of the confidence intervals mean that we cannot rule out the measure having an effect of less than 10 percentage points.

* *
*

Recruiter training is the main public measure used in France to combat discrimination in access to employment. Moreover, it was made compulsory for companies with more than 300 employees by the Equality and Citizenship Act of 29 January 2017. In this study, we assess the effect of measures which are similar in content and intensity to this compulsory training. But the results obtained from difference-in-differences estimation on experimental data do not confirm their effectiveness. Discrimination against the applicant of North African origin existed before the treatment and was not significantly reduced by the training measures four months after their implementation. This result suggests that the public policy consisting of imposing a compulsory training measure of low intensity with mainly legal content is not an adequate response to the challenge of ethno-racial discrimination in the labour market. It would be advisable either to bolster the intensity of the training measures, by increasing the duration or frequency of the sessions, or to change the nature of these measures so that they may effectively change behaviours, or else add complementary measures. This also suggests that other measures taken at the company level, such as centralisation of the human resources department, could be more effective than training (Berson *et al*., 2020) in reducing discrimination.

However, given the size of the sample and in particular the limited number of companies present in the treated group, we cannot be certain of detecting a small or even a major effect of these measures. This is the main limitation of this study. It would therefore be interesting to reproduce this type of assessment, if possible on a larger scale, to see whether an effect can be detected. Each test has a certain probability of detection of the effect, which may be low, especially if the actual effect is small. Moreover, difference-in-differences estimation reduces the statistical power compared with a mere comparison of proportions between two groups. Ideally, therefore, the companies which will receive the training should be selected randomly. That would also eliminate the risks of bias due to any unobserved characteristics varying over time between the groups.

The second limitation is due to the fact that the results are potentially influenced by the time allowed between execution of the training measures and the second series of tests. The five-month time horizon covered by this study could be too long to detect any very-short-term effect of the measures.

Lastly, this assessment only examines whether recruiters respond in the same way to requests for information on job opportunities depending on the presumed ethnic origin of an applicant. It cannot account for discrimination in the subsequent stages of the recruitment process. The identification of discrimination at the stage of requests for information is, however, widespread and used in other studies (Anne *et al.,* 2022 on the labour market; Bunel *et al*., 2021 and Le Gallo *et al*., 2020 on the housing market). ☐

BIBLIOGRAPHY

Agan, A. & Starr, S. (2018). Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment. *Quarterly Journal of Economics*, 133(1), 191–235. https://doi.org/10.1093/qje/qjx028

Anne, D., Chareyron, S., Leborgne, M., L'Horty, Y. & Petit, P. (2022). Discriminations dans l'accès à l'emploi : une exploration localisée en pays Avesnois. *Revue d'Économie Régionale & Urbaine* N° 3, 413–430. https://doi.org/10.3917/reru.223.0413

Arrow, K. J. (1974). *Discrimination in Labor Markets*. *Princeton University Press*, pp. 1–33. https://doi.org/10.1515/9781400867066-003

Baert, S., Cockx, B., Gheyle, N. & Vandamme, C. (2015). Is There Less Discrimination in Occupations Where Recruitment Is Difficult? *ILR Review*, 68(3), 467–500. https://doi.org/10.1177/0019793915570873

Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press.

Behaghel, L., Crépon, B. & Le Barbanchon, T. (2015). Unintended Effects of Anonymous Résumés. *American Economic Journal: Applied Economics*, 7(3), 1–27. https://doi.org/10.1257/app.20140185

Benedetto-Meyer, M. (2024). Quelles sont les caractéristiques de l'offre de formation à la non-discrimination à l'embauche ? Dares, *Document d'études* N° 274.
https://dares.travail-emploi.gouv.fr/sites/default/files/4db8760fd807009418f2a81daa144506/Dares_DE_offre%20formation_non-discrimination.pdf

Berson, C., Laouénan, M. & Valat, E. (2020). Outsourcing recruitment as a solution to prevent discrimination: A correspondence study. *European Association of Labour Economists, 31st Annual Conference, Uppsala Sweden*, 19–21 September 2019, 64, 101838. https://doi.org/10.1016/j.labeco.2020.101838

Bertrand, M., Chugh, D. & Mullainathan, S. (2005). Implicit Discrimination. *American Economic Review*, 95(2), 94–98. https://www.jstor.org/stable/4132797

Bertrand, M. & Mullainathan, S. (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review*, 94(4), 991–1013.
https://doi.org/10.1257/0002828042002561

Bunel, M., L'Horty, Y., du Parquet, L. & Petit, P. (2021). La discrimination liée à l'origine sur le marché du logement à Paris. *Revue française d'économie*, vol. XXXVI, N° 1, 185–219.
https://doi.org/10.3917/rfe.211.0185

Chareyron, S., Challe, L., L'Horty, Y. & Petit, P. (2022). Can subsidies paid directly to employers reduce residential discrimination in employment? An assessment based on serial field experiments. *Urban Studies*, 59(6), 1202–1218. https://doi.org/10.1177/00420980211006033

Chareyron, S., L'Horty, Y., Mbaye, S. & Petit, P. (2023). Reducing ethnic discrimination through formal warning: Evidence from two combined field experiments. *Regional Science and Urban Economics*, 98, 103850.
https://doi.org/10.1016/j.regsciurbeco.2022.103850

Chareyron, S., L'Horty, Y. & Petit, P. (2023). Discrimination in Access to Employment: The Combined Effects of Gender, Origin and Address. *Economie et Statistique / Economics and Statistics*, 541, 3–16.
https://doi.org/10.24187/ecostat.2023.541.2104

Deuchert, E. & Kauer, L. (2017). Subventions à l'embauche et handicap : les enseignements d'une expérience à petite échelle. *Revue internationale du Travail*, 156(2), 299–317. https://doi.org/10.1111/ilrf.12058

Edo, A. & Jacquemet, N. (2013). Discrimination à l'embauche selon l'origine et le genre : défiance indifférenciée ou ciblée sur certains groupes ? *Economie et Statistique*, 464-466, 155–172.
https://doi.org/10.3406/estat.2013.10235

Fang, A. H., Guess, A. M. & Humphreys, M. (2019). Can the Government Deter Discrimination? Evidence from a Randomized Intervention in New York City. *The Journal of Politics*, 81(1), 127–141.
https://doi.org/10.1086/700107

Krause, A., Rinne, U. & Zimmermann, K. F. (2012). Anonymous job applications in Europe. IZA, *Journal of European Labor Studies*, 1(1). https://doi.org/10.1186/2193-9012-1-5

Le Gallo, J., L'Horty, Y., Du Parquet, L. & Petit, P. (2020). Discrimination in Access to Housing: A Test on Urban Areas in Metropolitan France. *Economie et Statistique / Economics and Statistics*, 513, 27–45.
https://doi.org/10.24187/ecostat.2019.513.2004

Murchie, J., Pang, J. & Schwegman, D. J. (2021). Can information help Lakisha and Jamal find housing? Evidence from a low-cost online experiment of landlords. *Regional Science and Urban Economics*, 90, 103712.
https://doi.org/10.1016/j.regsciurbeco.2021.103712

Riach, P. A. & Rich, J. (2010). An Experimental Investigation of Age Discrimination in the English Labor Market. *Annals of Economics and Statistics*, 99/100, 169–185. https://doi.org/10.2307/41219164

APPENDIX

Table A1 – **Positive response rates before and after treatment (Chambéry)**

	Before treatment	After treatment	Difference	Difference-in-differences
Treated group				
French origin	44.93	39.13	−5.80	
North African origin	30.43	23.19	−7.24	
Difference	−14.5*	−15.94**	−1.44	
Control group				−0.49
French origin	35.41	33.49	−1.92	
North African origin	21.53	18.66	−2.87	
Difference	−13.88***	−14.83***	−0.95	

Note: *** p<0.01; ** p<0.05; * p<0.1. The p-values are obtained from equality of proportion tests. The last column of the table shows the difference in variation of discrimination between the treated group and the control group and between the two periods (−1.44−(−0.95)).
Source: Diademe testing, TEPP-CNRS.

Table A2 – **Positive response rates before and after treatment (PACA)**

	Before treatment	After treatment	Difference	Difference-in-differences
Treated group				
French origin	45.86	42.11	−3.75	
North African origin	36.84	28.57	−8.27	
Difference	−9.02	−13.54**	−4.52	
Control group				−13.59
French origin	32.75	24.69	−8.06**	
North African origin	19.14	20.15	1.01	
Difference	−13.61***	−4.54	9.07**	

Note: *** p<0.01; ** p<0.05; * p<0.1. The p-values are obtained from equality of proportion tests. The last column of the table shows the difference in variation of discrimination between the treated group and the control group and between the two periods (−4.52−9.07).
Source: Diademe testing, TEPP-CNRS.

Table A3 – **Effect of the training measures (without companies in the control group with more than 300 employees)**

			Chambéry		PACA	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
NAfr	−0.122*** (0.040)	−0.120*** (0.040)	−0.165** (0.074)	−0.164** (0.075)	−0.067 (0.055)	−0.079 (0.055)
T × NAfr	0.014 (0.059)	0.017 (0.060)	0.045 (0.115)	0.044 (0.114)	−0.042 (0.086)	−0.019 (0.085)
Post	0.169 (0.184)	−0.067 (0.228)	−0.250 (0.303)	−0.273 (0.364)	0.114 (0.265)	−0.944*** (0.342)
NAfr × Post	0.039 (0.055)	0.041 (0.055)	0.056 (0.088)	0.056 (0.088)	−0.035 (0.075)	−0.004 (0.076)
T × Post	0.011 (0.064)	−0.003 (0.063)	0.116 (0.115)	0.114 (0.114)	−0.051 (0.088)	−0.100 (0.089)
T × NAfr × Post	−0.029 (0.086)	−0.036 (0.086)	−0.069 (0.148)	−0.069 (0.148)	0.076 (0.114)	0.039 (0.113)
Date fixed effects	X	X	X	X	X	X
Company fixed effects	X	X	X	X	X	X
Order of sending		X		X		X
Observations	3,132	3,132	1,072	1,072	2,060	2,060
R^2	0.163	0.167	0.294	0.294	0.178	0.196

Note: *** p<0.01; ** p<0.05; * p<0.1. The coefficients estimated by the ordinary least squares method are presented in the table. The standard errors clustered at the company level are shown in brackets. The control group observations are weighted by the inverse of the number of control subjects selected for each company treated.
Source: Diademe testing, TEPP-CNRS.

Introduction – Evaluation of Public Policies
A Selection of Papers Presented at the 9th Annual Conference on Public Policy Evaluation, Hosted by the AFSE and the Directorate General of the Treasury

Étienne Lehmann*

That evaluating public policy is a democratic duty can no longer be in doubt. Indeed, in France it is a constitutional obligation: Article 15 of the 1789 Declaration of the Rights of Man and of the Citizen states that "society has the right to hold to account any public agent of its administration." Necessity aside, however, the issue of how to convincingly evaluate public policy remains a fraught question. The challenge is not simply to observe what happens when a new policy is put in place. The real task is to compare the observed situation with what would have happened if that policy (and only that policy) had never been introduced, the so-called "counterfactual" scenario. The complexity arises from the fact that this counterfactual scenario is, by definition, not observed. The evaluation of public policy thus poses formidable methodological challenges, since counterfactual scenarios require modelling, a practice which is inherently debatable and open to criticism.

Although the evaluation of public policy and its attendant methodological challenges are not limited to the sphere of economic policy, evaluating governments' economic policies has nonetheless become both a fertile field of economic research and an essential requirement for the administration of those same economic policies. With this in mind, the French Economic Association (AFSE, *Association Française de Science Économique*) and the Directorate General of the French Treasury (*Direction générale du Trésor*) joined forces in 2015 to co-host an annual conference on the evaluation of public policy. Every year the conference selects papers on various aspects of public policy evaluation, with submissions assessed against two key criteria: academic excellence, and relevance to the administration of economic policy in France.

In 2023, in agreement with the editors of the journal, the organisers encouraged the authors of the papers presented at this 9th conference to submit their contributions for publication in *Economie et Statistique / Economics and Statistics*. The two articles included in this dossier are the fruit of this partnership.

* *Université Paris Panthéon-Assas, CRED.* Correspondence: *etienne.lehmann@assas-universite.fr*

Translated from: "Introduction – L'évaluation des politiques publiques. Sélection d'articles présentés à la 9e conférence sur l'évaluation des politiques publiques organisée par l'AFSE et la Direction générale du Trésor".

Citation: Lehmann, E. (2025). Introduction – Evaluation of Public Policies. A Selection of Papers Presented at the 9th Annual Conference on Public Policy Evaluation, Hosted by the AFSE and the Directorate General of the Treasury. *Economie et Statistique / Economics and Statistics*, 547, 17–19. doi: 10.24187/ecostat.2025.547.2136

In the first of these articles, entitled “The Macroeconomic Effects of the Energy Price Cap: An Evaluation Conducted Using the ThreeME Multisectoral Model”, **Paul Malliet and Anissa Saumtally** set out to analyse the impact of the energy price cap introduced in France in response to soaring global energy prices linked to the rebound from the COVID-19 pandemic and the international sanctions imposed in the wake of the Russian invasion of Ukraine. To do so, they make use of a computable general equilibrium model known as ThreeME. The originality of this model resides partly in its multisectoral scope, but also in its capacity to provide a finely detailed representation of energy flows within the economy. Last but not least, ThreeME is a “neo-Keynesian” model wherein prices do not instantaneously adjust to maintain market equilibrium. All of these features allow for very detailed analysis of price trends and the repercussions of price shocks from one industry to another. The authors conclude that the price cap appears to have cushioned the impact of global shocks on domestic energy prices by an amount equivalent to 0.2% of GDP in 2022 and 0.4% in 2023, compared to a reference scenario with no price cap. However, the cost of the measure to the French national budget is estimated to have been somewhere in the region of 0.5% and 0.7% of GDP for 2022 and 2023 respectively.

In the second article, entitled “The Distance Between Occupations, and Post-training Professional Transitions for Jobseekers”, **Kevin Michael Frick, Yagan Hazard, Damien Mayaux and Thomas Zuber** set out to evaluate the extent to which the vocational training opportunities available to unemployed people actually succeed in attenuating the structural imbalances between the skills held by jobseekers and the skills required by occupations, especially for the occupations for which recruitment shortages are the most important. With this objective in mind, the authors adopted a methodology which we found to be particularly innovative. Their strategy is based on comparisons of real examples of occupational transition by jobseekers (who had previously been in stable employment in a different occupation), some of whom had received vocational training and some of whom had not. The authors sought to ascertain whether or not access to vocational training facilitates the redistribution of manpower towards sectors experiencing recruitment shortages.

In order to achieve this goal, they began by constructing an indicator capable of measuring the skill gap between two occupations. They then used this tool to measure, for each jobseeker in the sample, the distance between their previous occupation and the occupation which they took up when returning to employment. The methodological originality of their study owes much to their use of a machine learning technique involving a neural network. They began by training the model to characterize almost four million job offers posted on the Pôle Emploi website, through a vector identifying twenty core characteristics. The difference between two occupations could thus be represented in geometric terms, by the angular distance between the two vectors as well as the norm for those vectors. Using this measurement tool, the authors were able to estimate – for each individual case of a jobseeker who previously held a stable post and later returned to work in a new profession – the distance between the old and the new occupations.

In order to estimate the impact of vocational training, the authors compared the career trajectories of jobseekers who had undergone training and jobseekers who were similar in terms of other observable characteristics, but had not received training. Instead of resorting to a familiar method such as propensity score matching, the authors made use of a more flexible method called *Double Debiased Machine Learning*. This method enabled the authors to interpret their results as measurements of the correlation between vocational training and the return to employment of unemployed individuals, corrected for observable differences in other characteristics. The authors thus conclude that vocational training is correlated with a reduced probability of returning to work in one’s original occupation, and a greater probability of returning to employment in an occupation demanding very different skills from those associated with that original occupation.

Above and beyond the light it casts on the role of vocational training in getting unemployed people back to work in new occupations, we feel that this article illustrates the potential of new machine learning techniques to revolutionise the econometric dimensions of public policy evaluation.

Taken together, these two articles illustrate the diversity of approaches which public policy evaluation can take, and, indeed, the diversity of the policies we seek to evaluate. Public policy evaluation remains a remarkably fecund field of research, and we look forward to welcoming you to the next Annual Conference on Public Policy Evaluation, co-hosted by the French Economic Association (AFSE) and the Directorate General of the Treasury. □

The Macroeconomic Effects of the Energy Price Cap: An Evaluation Conducted Using the ThreeME Multisectoral Model

Paul Malliet* and Anissa Saumtally*

Abstract – The energy crisis that struck Europe in 2021 as the world bounced back from COVID, and amplified by the Russian invasion of Ukraine, led to a sharp increase in energy prices, particularly gas prices. In this context, European nations implemented emergency measures to protect households' purchasing power and the competitiveness of their businesses. France chose to mitigate energy price rises by implementing a price cap. Making use of a computable general equilibrium model, we explicitly simulate the divergent trajectories of energy prices with and without this price cap. Our results show that the budgetary cost of this measure was lower than initially expected, and while the macroeconomic impact was also relatively small, it did nonetheless preserve household purchasing power.

JEL: C68, E64, E65, Q43, Q48
Keywords: macroeconomics, energy crisis, energy price cap, public policy evaluation

**Observatoire Français des Conjonctures Économiques, Sciences Po.* Correspondence*: anissa.saumtally@sciencespo.fr*

The authors would like to thank Xavier Ragot for the opportunity to publish this work, which was originally published as a Focus bulletin by the CAE. The authors would also like to thank the members of the CAE, the participants in the OFCE internal seminar, the participants in the 9th AFSE-DG Trésor (Directorate-General for the Treasury) conference on the evaluation of public policies, and the two anonymous reviewers for their comments and remarks. The authors would also like to thank the ADEME for financially supporting the development of the ThreeME model under the aegis of research convention 23ESD0184.

Received in March 2024, accepted in November 2024. Translated from "Les effets macroéconomiques associés au bouclier tarifaire : une évaluation conduite à l'aide du modèle multisectoriel ThreeME".
The opinions and analyses presented in this article are those of the author(s) and do not necessarily reflect their institutions' or INSEE's views.

Citation: Malliet, P. & Saumtally, A. (2025). The Macroeconomic Effects of the Energy Price Cap: An Evaluation Conducted Using the ThreeME Multisectoral Model. *Economie et Statistique / Economics and Statistics*, 547, 21–47 (First published online: September 2025). doi: 10.24187/ecostat.2025.547.2137

The Russian invasion of Ukraine in February 2022 exacerbated a major energy crisis for the European Union, which had first emerged in September 2021 in the wake of the post-COVID rebound in international demand. Although European countries were quick to condemn Russia and introduce economic sanctions in February 2022, particularly on energy products such as coal and oil,[1] their heavy dependency on Russian gas for energy supplies[2] posed a major risk to the stability of their energy networks and the continued functioning of their economies.

In spite of the uncertainties expressed in 2022 with regard to the capacity of Europe's energy system to withstand the sudden withdrawal of Russian energy imports, the system has in fact demonstrated its resilience. As noted by the International Energy Agency (IEA, 2023), various factors contributed to a 3% reduction in global energy consumption in 2022 compared with the previous year. These included a mild winter in Europe, along with conscious efforts to reduce energy consumption.

Another important aspect of the crisis is the way in which it has amplified the operating fluctuations of European electricity markets, particularly for intraday and day-ahead transactions. On the futures market, exchanges between energy buyers and producers are agreed in advance and the prices and quantities are fixed. Such transactions generally involve power plants which can be controlled, and the prices recorded are generally lower than those seen on the day-ahead and intraday markets. The latter markets are constantly balancing supply against demand, and sale and purchase agreements are made for fixed periods of time. The equilibrium price is also known as the spot price.[3] It corresponds to the marginal cost of the most recent production unit put into service, following the *merit order* principle (power stations are utilised in a specific order based on their respective production costs: the cheapest sources are prioritised, then progressively more expensive sources are used until the demand is satisfied), with all producers being paid at this marginal price. To the extent that gas-fired power plants (and, to a lesser extent, fuel oil plants) can be controlled, and thus serve to guarantee the stability of the network, the spot prices for gas are often used as the benchmark for setting electricity prices, with the Dutch TTF regarded as the reference market (Figure I).

Between December 2020 and December 2021, the price of importing energy into the

1. A complete ban on imports of petroleum refining products was introduced in February 2023.
2. It still accounted for half of its imports on the eve of the invasion of Ukraine, and from 2023 onwards it accounts for less than 10% of its total imports (source: Bruegel, https://www.bruegel.org/dataset/european-natural-gas-imports).
3. The difference between the cost of production and the price paid is known as the infra-marginal rent.

Figure I – **Variations in the Dutch TTF index for gas spot prices**

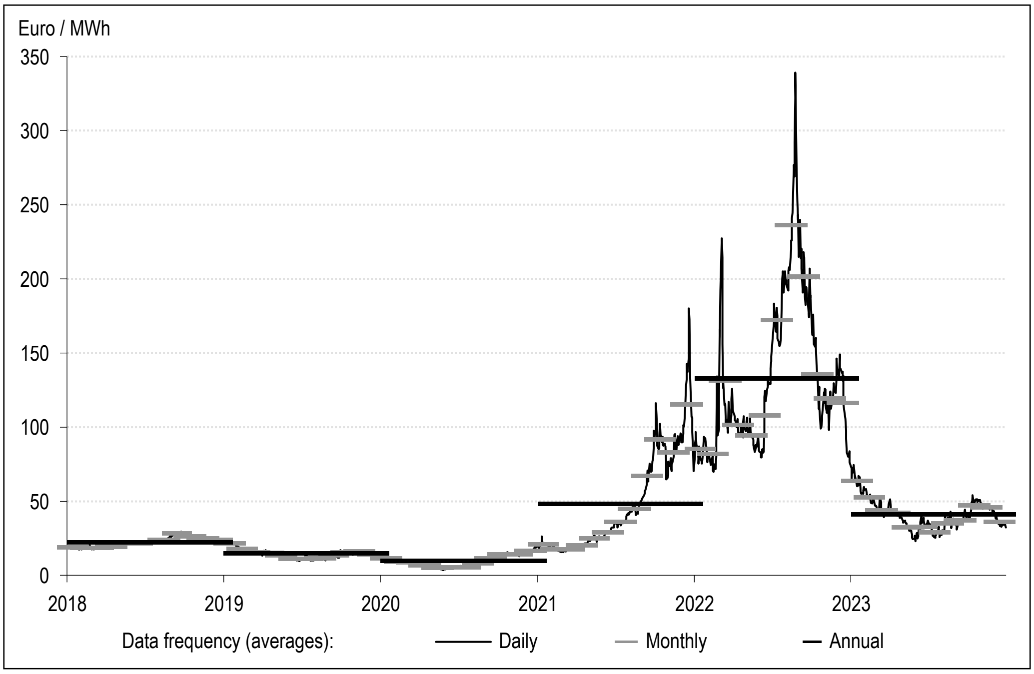

Note: The horizontal lines represent the annual (black) and monthly (light grey) averages.
Source: ICE.

Eurozone more than doubled, driving inflation up in European nations. Average inflation in EU member states was 9.2% in 2022, a threefold increase on the preceding year. France was an exception to the rule, with a rate of inflation of 5.9%, while the rate rose to 8.3% in Spain, 8.7% in Germany and Italy, 11.6% in the Netherlands and 13.2% in Poland.

With regard to the situation in France, between Q2 2021 and Q2 2022, rising energy prices contributed 3.1 percentage points (pp) to the total rate of inflation in France, which stood at 5.3%. The introduction of a *price cap* ("*Bouclier tarifaire*" in French) limiting price rises for electricity and gas to 4% made it possible to reduce energy inflation from 54.2% to 28.5% for households (and 50.3% to 20.3% for businesses) (Bourgeois & Lafrogne-Joussier, 2022). The energy price cap operates on the principle that the national government will subsidise the difference between the capped price paid by consumers and the price charged by suppliers, determined by market conditions. As such, while it maintains inflation at the desired level, the cost of this measure for the government depends primarily on market prices. Among the measures adopted by European nations, France's decision to maintain prices at a certain level was relatively unusual (Sgaravatti *et al.*, 2023); most other EU member states opted for transfer mechanisms.

The purpose of this study is to better understand the macroeconomic effects of the energy price cap by explicitly representing the price structure, from wholesale prices to consumer prices, within the macroeconomic context. We conducted this evaluation with the help of a computable general equilibrium model called ThreeME, combined with a detailed calibration of electricity and gas prices. An alternative approach to evaluating the macroeconomic effects of the energy price cap in France (Langot *et al.,* 2023) has been proposed, based on a *Heterogenous Agents Neo Keynesian* (HANK) model. The authors of that study estimate that the price cap served to reduce inflation by 1.1 pp in 2022 and 1.8 pp in 2023, while mitigating the decline in GDP growth by 1.1 pp to keep it at 2.9 pp in 2022, and by 0.9 pp to keep it at 1pp in 2023, all at a fiscal cost of approximately 2% of GDP.

These two approaches reveal themselves to be more complementary than contradictory, to the extent that they use different models with different theoretical frameworks and refinement techniques, thus providing multiple perspectives on the same issue. The HANK model, for example, provides an integrated representation of household heterogeneity, thus enabling us to assess the redistributive dimension of the policy, a possibility not offered by the ThreeME model since it focuses on a single representative household. On the other hand, using a detailed multisectoral model enables us to establish an explicit representation of price trends for different energy products (gas and electricity). In the calculations made using the HANK model,[4] the price cap is regarded as an additional spike in public spending with an impact on an economy comprising both a composite product and an energy product: it is based on the forecasts issued by the French government in the Draft Budget Bill for 2023,[5] presented to the Parliament in September 2022 and passed into law in December 2022. One of the limitations inherent to this approach is that it relies upon the government's estimates for the expected fiscal cost of the measure, estimates which were calculated when the spot price was at its peak (cf. Figure I) and thus fail to reflect the waning of prices observed post-26 August 2022. Our study seeks to further explore this matter, incorporating an updated price estimate into the macroeconomic framework.

Section 1 offers some context on the French energy market and the data used to calibrate the model. In Section 2, we detail the modelling framework we used to analyse the energy price cap policy. In Section 3 we present our results, which are further discussed in Section 4, before offering a brief conclusion.

1. Political Context

1.1. The French System: A Two-Speed Energy Market

Formerly a state-controlled monopoly, the business of supplying gas and electricity was opened up to new actors in 2007, with a view to creating a competitive marketplace. Although the market is now open, electricity consumers in France can still choose between market prices and regulated prices. This was also true of gas until July 2023, when the two-speed system came to an end. The regulated price scheme available to consumers is operated by France's historic energy supplier, with prices set by the Energy Regulation Commission *(Commission de régulation de l'énergie,* or CRE), an independent body not under government control. The regulated

4. To assess the relevance of these models in the study of energy shocks, see Auclert et al. *(2023).*

5. See the documentation associated with the 2023 Draft Budget Bill (called PLF *in French) : https://www.budget.gouv.fr/documentation/documents-budgetaires/exercice-2023/le-projet-de-loi-de-finances-et-les-documents-annexes-pour-2023* energy price comprises three components: fair remuneration for the energy supplier (dependent upon wholesale prices), distribution and network costs, and taxes (VAT and other forms of excise). Each of these components accounts for around a third of the price. Prices are set annually for electricity, and monthly for gas. For electricity prices, this means that consumers are guaranteed a stable price for the entire year. For suppliers, if wholesale prices vary considerably over the course of year, to the extent that the regulated price is no longer enough to cover their costs, then a price supplement will be added for the subsequent period so that they may recoup those costs. The more frequent adjustment of gas prices, meanwhile, creates greater volatility for consumers, but allows for better adjustment to fluctuations in wholesale prices. Due to the frequency with which prices are adjusted, the abolition of regulated gas prices in July 2023 has not significantly impacted the gas price trend.

The gas price deals available on the market are generally based upon the observed regulated prices. Often, households sign up to contracts which guarantee them gas at a price several percent below the regulated price for a fixed period of time, after which time the price is recalculated with reference to current market conditions. Alternative suppliers cover their costs by anticipating wholesale prices on the futures market, and optimising their contract prices. As such, at any given moment, average prices should be more or less aligned with the regulated price. Due to the unexpected spike in wholesale prices, some contracts turned out to be unprofitable because they were locked into low prices. Conversely, new contracts signed at market prices during the price surge would have involved significant price increases. Households signed up to market price contracts currently represent approximately 30% of the overall market; the vast majority of households prefer to stick with regulated price contracts.

On account of this system, consumer prices for energy in France usually see relatively modest fluctuations (Figure II), because energy prices tend to be anchored by the regulated price. The repercussions of variations in wholesale energy prices in France are among the weakest in Europe, particularly for electricity (Ari *et al.,* 2022). The same goes for the contribution of energy prices to headline inflation in France, compared with other European nations.

1.2. Regulating Energy Prices in Times of Crisis: The Price Cap

In the second half of 2021, wholesale gas prices saw a series of major spikes (fluctuations in the average daily price were as much as ten times greater than pre-crisis levels), having lingered at historically low levels in 2020 in the midst of the COVID-19 pandemic. As economies rebounded rapidly from the crisis, especially in Asia, and Europe experienced a cold winter, the demand pressure for natural gas led to a first significant increase in wholesale prices. This

Figure II – **Monthly indices for gas and electricity prices**

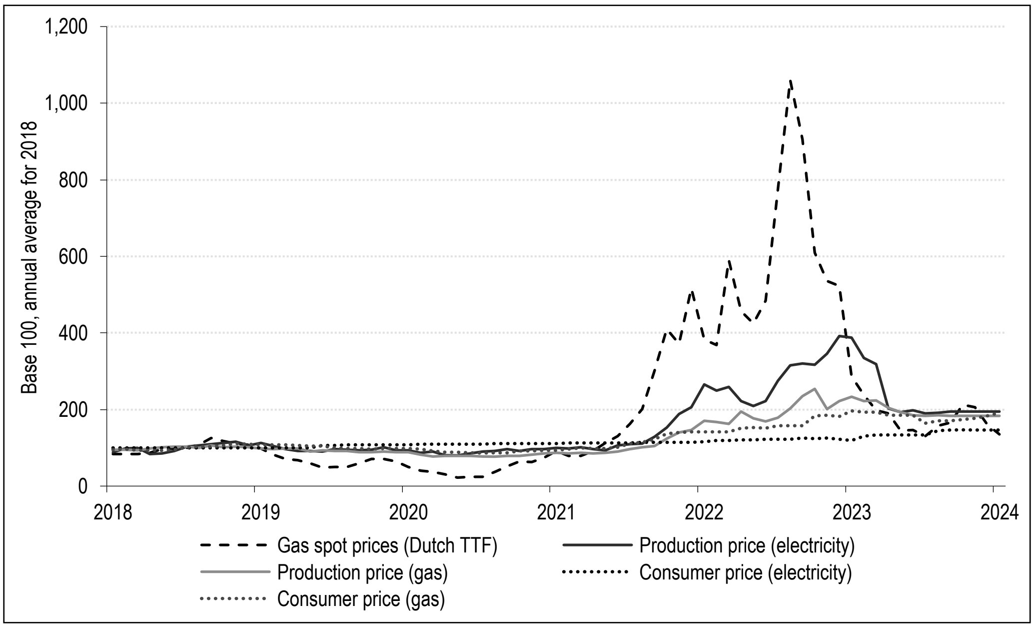

Note: The dotted lines correspond to the annual averages.
Sources: ICE, INSEE.

energy inflation crisis was exacerbated by the second Russian invasion of Ukraine, starting 24 February 2022. When the European Union declared an embargo on Russian gas, with great difficulty on account of the severe gas dependency of some member states, the cost of alternatives (primarily liquefied petroleum gas, LPG) increased further still. On the Dutch natural gas futures market (Dutch TTF), the benchmark for wholesale gas and electricity prices in Europe, intraday prices exceeded the 1,000 euro/MWh threshold on several occasions.

In France, where regulated gas prices are recalculated monthly, the method employed by the CRE to calculate these prices would normally have integrated these fluctuations, leading consumer prices for gas to rise to levels which might have become prohibitively high for households. This marked the point at which the government began to intervene to steady gas prices. The first price capping measure consisted in freezing the regulated gas price for households at the end of 2021 and for the whole of 2022. For electricity, a similar mechanism was applied to the regulated price in 2022, keeping price increases in 2022 to 4% instead of the estimated 30% which had been forecast. In 2023, the price increase was capped at 15% for both domestic gas[6] bills and electricity. Without this measure, prices would have doubled.

In order to mitigate the impact of this price cut on suppliers' costs, the cap was partly made possible by removing some of the taxes usually paid by electricity consumers (the Interior Tax on End Electricity Consumers; TICFE – and the Local Tax on End Electricity Consumers; TCCFE). Supplementing this measure, the CRE announces a theoretical price – the price which would have been applied where there no price cap – and the government pays the difference directly to the suppliers. To assess the total cost of this policy, we must add up the value of the tax discount and the subsidy as shown in Figure III below, which shows the resulting mean theoretical and applicable prices; in this graph, the difference between the dotted lines and the solid lines represents the cost of the measure (excluding the tax discount, because the prices shown are pre-tax) to the government.

Regulation of gas prices came to an end on 30 June 2023, at the same time that the price cap scheme for gas was wound down. Since that date, the CRE has published a "guide price" calculated in much the same way as the old, regulated retail price (hereafter "RRP"), although suppliers are free to set their own retail prices. In Figure III, and elsewhere in this study, we use

6. *For gas, this price increase is limited to the first half of the year. Given the fall in wholesale gas prices, the gas price cap has not been extended to the rest of the year.*

Figure III – **CRE regulated retail prices for gas and electricity**

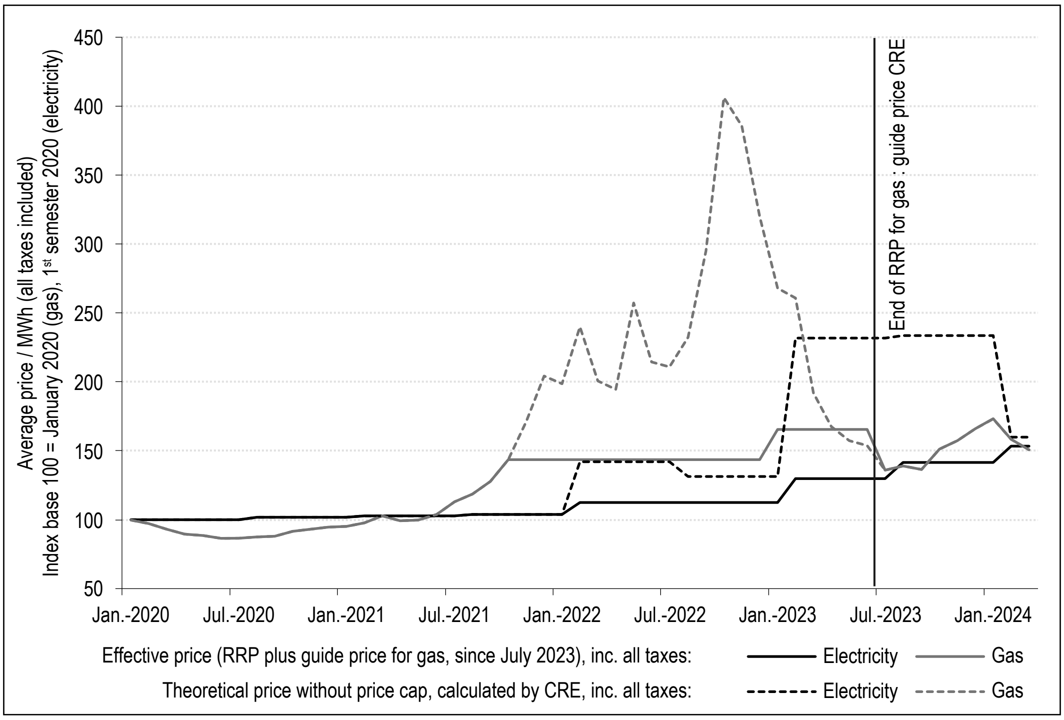

Note: The dotted lines correspond to the annual averages.
Source: Energy regulation commission, authors' calculations.

this guide price as the price indicator for June 2023 onwards (see Figure A3-II in Appendix 3).

As for electricity, in light of falling gas prices the government decided that the price cap would come to an end in 2025, and from 2024 onwards the TICFE would be gradually increased from the 0.0001 euro/kWh rate in place since the capping policy was introduced (in late 2021 it was around 0.03 euro/kWh). In 2024, the regulated retail price for electricity (RRPe) published by the CRE once again became the benchmark price, i.e. it was considered sufficient to pay energy suppliers without any further compensatory remuneration from the government; this measure has been wholly funded by an initial tax increase (see Figure A3-I).

In practice, the aim of this measure is to cap the unit price of a kilowatt-hour of energy, including all taxes. In theory, there are no conditions pertaining to the quantity of energy consumed. However, in the meantime the government also launched a campaign encouraging consumers to reduce their energy consumption, raising the spectre of electricity shortages if demand were to outstrip supply capacities.[7] In 2022, electricity and gas consumption did fall in spite of what was a relatively small increase in prices. Nonetheless, it is difficult to determine whether this was a direct consequence of campaigns encouraging more responsible energy consumption, or else a trend motivated by fears of further price rises.

The most recent figures published by the French government in the Stability Programme (PSTAB) for 2024[8] (published in April 2024) estimate that this measure alone cost 4.5 billion euros for gas and 16.6 billion for electricity in 2022. In 2023, with wholesale prices falling sharply (by May 2023 they were back below their 2021 average), the cost of the policy was estimated at 24.3 billion euros for electricity and 2 billion euros for gas (see Figure IV). It should be noted that the high estimated cost of the electricity price cap is primarily a result of the design of the pricing system, which incorporates a certain latency into the RRP-setting mechanism.

As mentioned above and demonstrated in Figure III, regulated electricity prices are calculated annually by the CRE, with minor adjustments in the second half of the year. The majority of the theoretical price increase (which determines the cost of this measure) can probably be attributed to the remuneration paid to suppliers, including compensation for losses sustained in the previous year due to the unforeseen increase in market prices. A subsequent update[9]

7. At the time, electricity production in France was slowed down by the temporary closure of certain nuclear power stations.
8. See the Stability Programme 2024-2027: https://www.tresor.economie.gouv.fr/Articles/2024/04/17/article-presentation-du-programme-de-stabilite-2024-2027 published in April 2024 by the Treasury Department.
9. See the update of 5 June 2023 at the following address: https://www.cre.fr/Actualites/la-cre-communique-des-clarifications-sur-les-dispositifs-de-boucliers-electricite-et-gaz-et-d-amortisseurs-electricite-pour-l-annee-2023-et-les-fo.html

Figure IV – **Estimating the cost of the price cap policy in successive annual budgets**

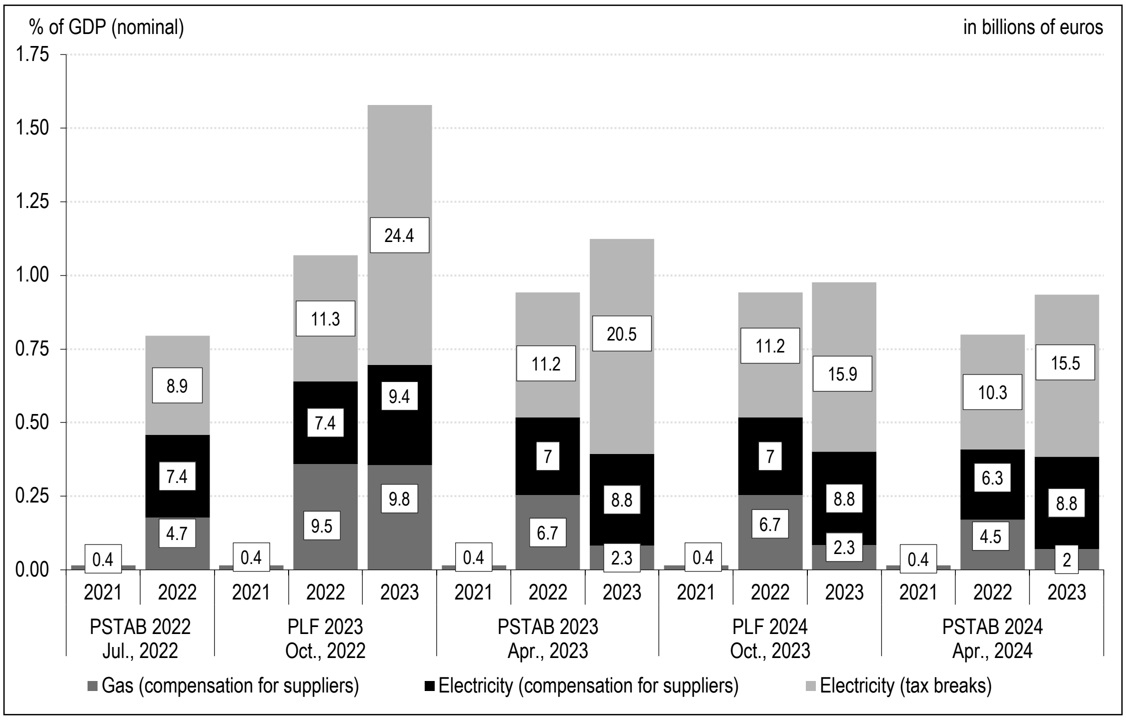

Source: The French Treasury.

from the CRE explained that the theoretical price of gas was capped at the level dictated by the energy price policy, while in fact it could have been lower on account of the current state of gas wholesale prices.

Furthermore, it is worth noting that the government's own estimates of the cost of these measures have changed significantly from one budgetary exercise to the next (Figure IV), testament to the complexity of budget forecasting in an uncertain context defined by an unprecedented crisis.

Efforts to keep consumer prices down also included a number of targeted measures such as direct subsidies for low-income households ("energy cheques"), measures which cost relatively little compared with the price cap policy. In this article, we omit to consider the increase seen in fuel prices, although it should be borne in mind that a temporary reduction in fuel prices, subsidised by the government, was also introduced in 2022 and benefited all consumers, regardless of the quantities they consumed. According to the government estimates published in the Stability Programme for 2024, measures taken to mitigate inflation are believed to have cost 39.5 billion euros in 2022 and 33.9 billion euros in 2023, with the price caps on gas and electricity accounting for 54% and 78% of those costs respectively.

2. Modelling Framework and Scenarios

2.1. The ThreeME Model

ThreeME is a computable general equilibrium (CGE) small open economy model,[10] originally developed to help decision-makers design and evaluate measures for decarbonising the French economy (Callonnec *et al.*, 2013; Hamdi-Cherif *et al.*, 2022; Callonnec & Cancé, 2022). ThreeME is specifically designed to evaluate the short, medium and long-term impacts of energy and environment policies at the sectoral and macroeconomic levels. To this end, the model combines several key characteristics:

- Its sectoral disaggregation allows us to analyse activity transfers from one sector to another, particularly in terms of employment, investment, energy consumption and trade balance;
- Its very detailed representation of energy flows within the economy enables us to analyse the consumption behaviour of economic agents regarding energy;
- Sectors may choose between capital and energy when relative energy prices rise, turning to substitute energy vectors;
- Consumers may make substitution decisions between energy vectors, modes of transport or consumer goods.

As a CGE model, ThreeME takes account of feedback effects between supply and demand (Figure V), as demand (consumption and investment) informs supply (production). Symmetrically, supply can stoke demand by means of the income generated by factors of production (labour, capital, energy products and materials). Compared with *bottom-up* energy models such as MARKAL (Fishbone & Abilock, 1981) or TIMES (Loulou *et al.,* 2005), ThreeME goes beyond simply describing the sectoral and technological dimensions of energy, instead integrating them into a comprehensive macroeconomic model.

ThreeME is a neo-Keynesian model, whereas existing CGE models in the Walras tradition largely concern themselves with supply, and prices are not adjusted instantaneously to achieve market equilibrium. This model is dynamic, and prices and quantities evolve slowly as producers adjust their supply to meet demand. One benefit of this is that it allows for short and medium-term periods of disequilibrium in the market (particularly periods of involuntary unemployment), creating a framework which is particularly conducive to the analysis of economic and energy policies.

This maximises the utility of each agent in period t subject to various constraints, such as market equilibrium (e.g. demand being equal to supply). This is a recursive-dynamic (i.e. short-sighted) model, which means that it begins by optimising the period t and then uses the endogenous results (for example, prices, wages and production levels) to optimise the ensuing period (i.e. $t+1$). Once the model has optimised the final period (as determined by the user), it generates forecasts for endogenous parameters such as prices, household income, GDP and the employment rate, for the whole time frame. ThreeME also requires a number of exogenous parameters: the *Social Accounting Matrix* (SAM) for the reference year, along with forecasts for population growth, the productivity of factors, and various substitution elasticities. SAM is a comprehensive database for the national economy, recording all transactions between economic agents at a given date (Kehoe, 1996). Forecasts for population and economic growth determine the availability of manpower,

10. The model is open source and its code is available at the following page: https://github.com/ThreeME-org/ThreeME_V3-open

Figure V – **Schematic representation of the ThreeME model**

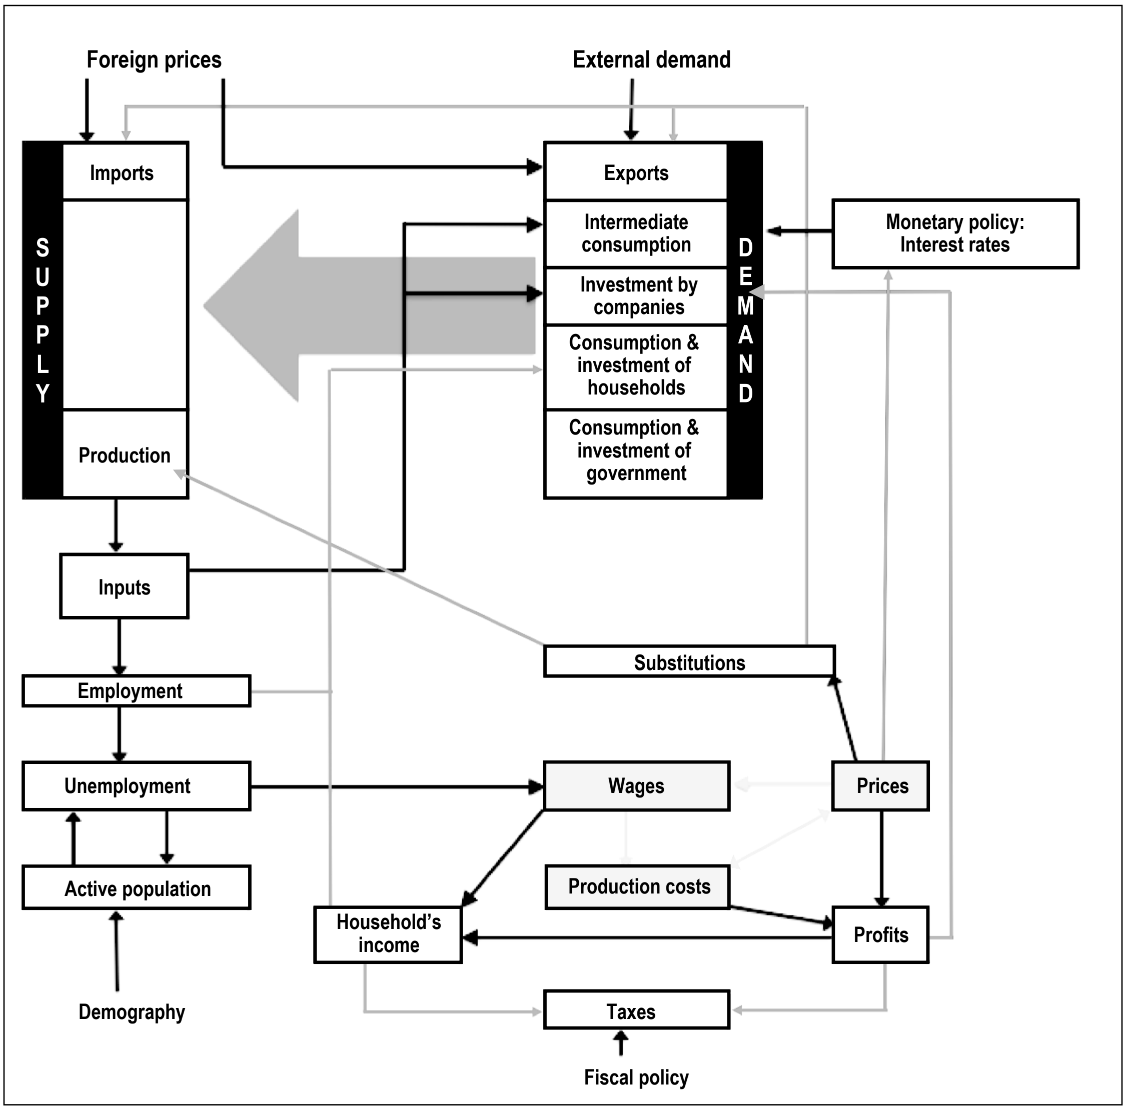

Source: https://www.threeme.org/documentation

and shape productivity trends. Elasticities determine the degree of substitution between factors of production within functions of production. With a *Constant Elasticity of Substitution* (CES) function, substitution between factors of production may follow either a linear production function, a fixed proportions (or Leontief) production function, or else a Cobb-Douglas production function. A linear production function represents a production process in which the factors of production are perfect substitutes (for example, labour can be completely replaced by capital). A fixed proportions production function represents a production process in which the factors of production are needed in fixed proportions. In the Cobb-Douglas production function, inputs may be substituted, even if they are not perfect substitutes. ThreeME uses a nested CES production function (Reynès, 2019) to describe the substitution between factors of production. This CES production function, known as KLEM, combines four factors of production: capital (K), labour (L), energy (E) and materials (M). Factors of production may be substituted for one another, with elasticity of substitution parameters determining the degree of substitution between them. Each pairing (i.e. $K-E$, $KE-L$, $KEL-M$) has its own substitution elasticity, explained in greater detail in the description of the model (Reynès *et al.,* 2021). One essential characteristic of standard, Neo-Keynesian AS-AD (aggregated supply and demand) macroeconomic models is that demand determines supply. Demand includes consumption (intermediate and final), investment and exports, while supply comes from imports and domestic production. By means of various feedback mechanisms, potentially involving some delay, supply shapes demand. The level of production determines the quantity of inputs used by businesses, and thus the quantity of their intermediate consumption and investment, two major components of demand. It also determines the level of employment and, as a result, influences final household consumption. Another effect of employment on demand is its influence on wages, by means of the unemployment rate which also depends upon the size of the active population. The size of the active population is primarily determined by exogenous factors such as demographic trends, but it is also shaped by endogenous factors such as the labour market participation rate.

2.2. Calibrating the ThreeME Model and Integrating CRE Data

For the French context, the ThreeME model was calibrated using data from the national accounts, available from Eurostat. Our reference year was 2015. After that reference year, the only shock we took into consideration was the global increase in energy prices beginning in 2021, in order to represent and isolate the trends observed in energy prices and analyse them independently of any other economic fluctuations. This shock was accounted for by integrating the increase in the Dutch TTF with a transmission coefficient of 50%.[11] In order to model consumer prices of energy in France, we made a further adjustment to the consumer price setting mechanisms in order to accurately reflect the price regulation structures described above. All of the other equations retain the standard specifications of the ThreeME model as usually applied in France.

2.2.1. Energy Prices

Within the framework of this model, the price equations rely on adjustment processes (see Appendix 2). For the purposes of this article, we modified the household price equations for two energy products: electricity and gas. In order to integrate the effects of the energy price cap policy, the two pricing equations for both energy products are determined exogenously so as to reflect the administered nature of these prices, in both the reference scenario and the scenario integrating the price cap as the price observed during the preceding period, to which we apply the price increase ratio fixed by the government (in both scenarios).

In formal terms, this can be written as follows:

$$P_{ce,t}^{CH} = \left(1 + \tau_{ce,t}\right) P_{ce,t-1}^{CH} \qquad (1)$$

where $P_{ce,t}^{CH}$ is the consumer price index for the period t for the energy products ce, gas and electricity, and $\tau_{t,ce}$ the annual growth rate calculated exogenously. For other consumer goods, we used the default pricing mechanism associated with this model, which submits prices to a process of adjustment. In order to model the impact of government policy, we include a reference price which corresponds to the CRE's theoretical price, alongside the regulated price. The difference between these two prices represents the cost of the policy to the government. In our reference scenario the two prices are equal, so the cost of the policy is essentially zero. Unlike the real policy framework surrounding the price cap, whose tools are broader in their scope, our simulations do not include that part of the policy financed by the tax exemptions mentioned in Section 2.2, simply because the taxes in question do not have a return effect for consumers within the standard ThreeME model. This means that the results of our simulations concerning the cost of the policy cannot be directly compared with the observed budgetary cost.

In order to obtain the energy price series used to calibrate our scenario, we calculated average consumer prices using the RRP. For electricity, the CRE publishes RRPe figures applicable under the price cap scheme (with and without taxes), as well as theoretical RRP figures without taxes. The theoretical RRP including tax was thus recalculated, based on the assumption that without the price cap the TIFCE tax would have remained at its late-2021 level. On the basis of these two data series, we used the standard consumer profile defined by the CRE, signed up to the "blue" basic price scheme[12] with an annual energy consumption of 2,400 kWh, at a power of 6 kVA. This enabled us to estimate an average price-per-kWh, taking fixed costs into account.

A similar approach was adopted for gas, where the CRE standard customer used for the purposes of our study is a consumer using gas for hot water, cooking and heating, with an annual consumption of 13.48 MWh at local price level NP2.[13]

The trend evolution of energy product prices, as modified in ThreeME for the purposes of this study, is presented in Figure VI hereunder. With the exception of the period 2021-2024, the trend for all prices was a year-on-year increase of +2%. For wholesale prices, based on empirical

11. This coefficient is taken from the study by Hernnäs et al. *(2023). It was estimated between the Dutch TTF and the selling price on the gas market during the period between April 2021 and August 2022 and for all European Union Member States.*
12. The RRPe determined by the CRE is defined for regulated 'basic' tariffs, but also for peak/off-peak contracts and, more recently, for 'Tempo' contracts.
13. This is the price level applicable to the largest number of consumers in the CRE's internal classification.

observations of the Dutch TTF, we recorded a price shock consisting of a 162% increase in 2021, repeated in 2022. We then reduced prices by 43% in 2023 and 2024. The shock applied until 2021 (inclusive) differs from the empirical fluctuations of the Dutch TTF, because the annual average price of gas hit a historic low in 2020, before increasing by 400% in 2021 and 177% in 2022. For the ThreeME simulations, we decided to smooth the variation in gas prices between the reference year (2015) and 2022 in order to obtain a stable trend for the years leading up to 2022, so as to facilitate both the calculations and the analysis of the results.

The reference scenario used for comparison corresponds to a world where no effort is made to keep consumer prices down. As such, we took the theoretical average increase in prices (before tax) determined by the CRE as our point of reference. For the purposes of our analysis, we then modified this consumer price specification to include the pre-tax price fixed under the price cap scheme.

We made one final modification to the model to ensure that energy consumption responded quantitatively to changes in prices, as seen in 2022 (corrected for meteorological effects), with

Figure VI – **Example of energy prices calibration in the ThreeME model**

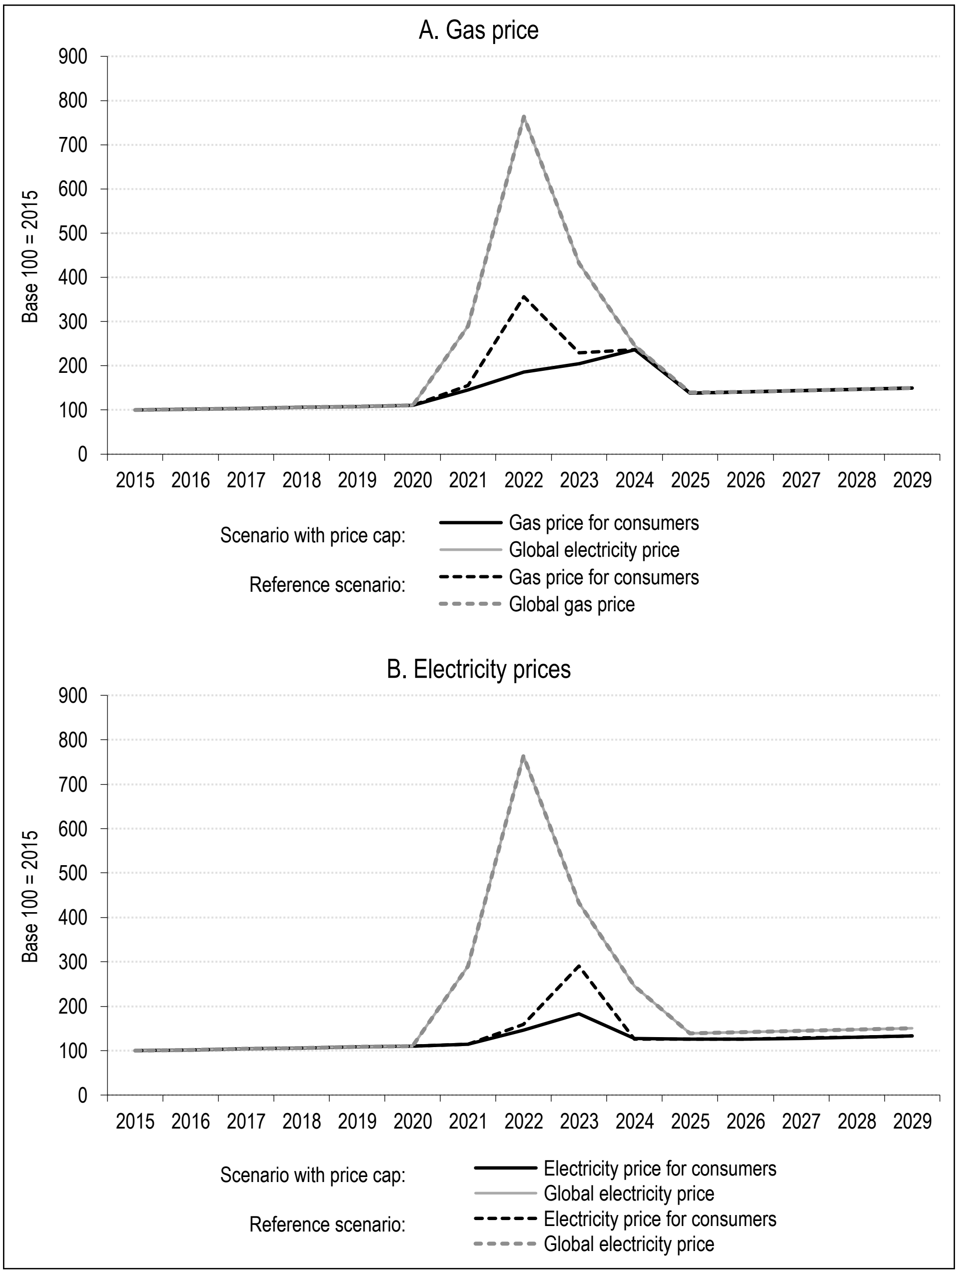

Source: ThreeME simulations.

electricity and gas consumption falling by 1.7% and 6.2% respectively.[14] We also calibrated the proportion of household energy demand which can be regarded as autonomous consumption in order to replicate the imputed decline in consumption for the years 2022 and 2023.

3. Results

We ran simulations using the ThreeME model for a period of 35 years, starting from our reference year 2015 (the year for which the model was calibrated using Eurostat data) and modelling the two scenarios described above.

Throughout the rest of our analysis we compare the results of the price cap policy with a reference scenario, which is essentially a variant where a shock affecting global energy prices upsets the steady state ThreeME model, along with an exogenously-determined variation in the consumer prices paid by households, reflecting the energy pricing policy adopted in France. The steady state scenario, meanwhile, is constructed with a steady growth rate of 1.25% and stable inflation of 2%. Figure VII presents a comparative plot of the price shock scenarios and the steady state scenario, giving a clearer picture of the impact of the crisis which is modelled hereunder. This graph shows that, at its peak in 2023, the price shock caused GDP to fall by around 0.4%, primarily because household consumption fell by more than 1.6%. We also observed an increase in investments, in spite of the slowdown in activity which was expected during the crisis. This can be attributed to substitution between energy and capital as factors of production, in response to rising energy prices. The GDP increase seen in 2021 can be attributed to falling gas imports, with the effect of the electricity price shock not felt until 2022.

The third scenario, the "price cap scenario", incorporates the price capping measures imposed on gas and electricity sales to private consumers. It is important to bear in mind that our calculations do not include any other price shocks, such as the spike in inflation observed from the second half of 2022 onwards, nor did we seek to model the economic impact of COVID-19. Our goal was to isolate the energy price shock and the political response. As a result, the majority of the results of this simulation are best understood in terms of the relative difference between the alternative scenario and the reference scenario.

We can see that incorporating the price cap increases real GDP by 0.2% in 2022 and 0.4% in 2023, as demonstrated in Figure VIII. This increase can be primarily attributed to household consumption, which is 0.8% greater in 2022 than it would have been had no measures been taken to mitigate energy price inflation. This indicates

14. Electricity consumption is set at 414 TWh (compared to 435 TWh in 2021) and gas consumption at 431 TWh (compared to 471 TWh in 2021).

Figure VII – **Price shock scenarios, compared with ThreeME stationary scenario**

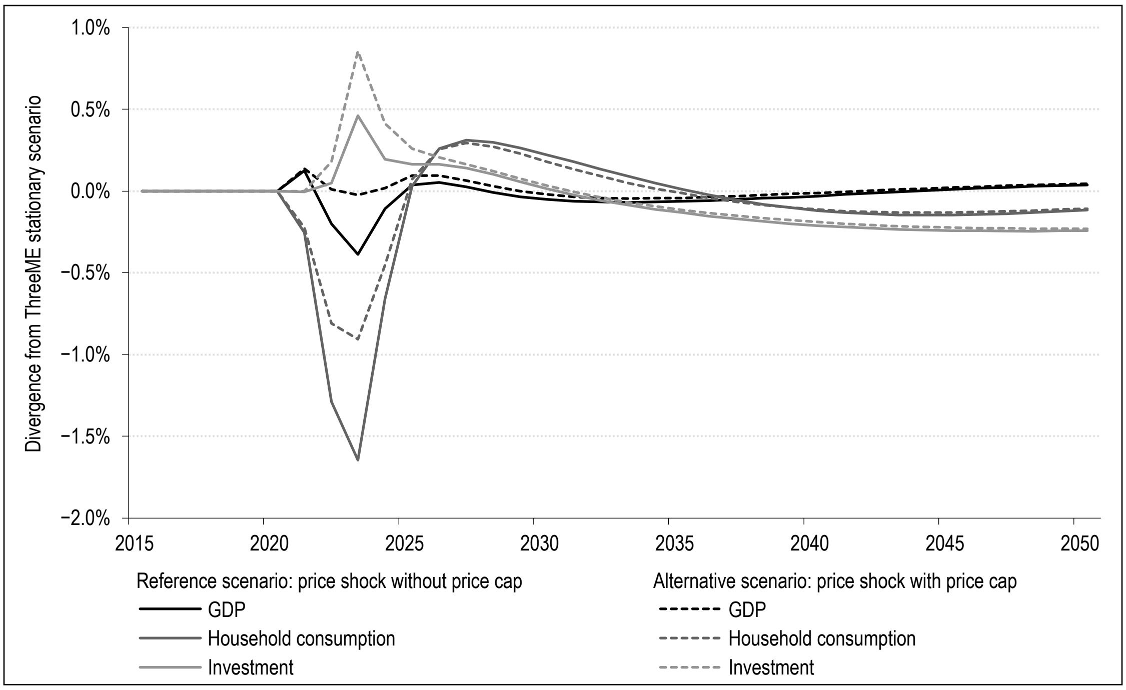

Note: The variables are given in volume terms.
Source: ThreeME simulations.

that the policy has indeed protected consumers' purchasing power, as illustrated by the difference in energy consumption.

Completing our analysis of these results, Figure VII indicates that the price cap scheme nearly halves the decline in household consumption in 2023, and virtually cancels out the decline in activity caused by the price shock in 2022 and 2023.

Consumption of natural gas falls by 5.2% in 2022, instead of the projected 17.3% (see Figure A1-I in Appendix 1). Similar results are obtained for electricity consumption in 2023, which falls by 5.9% instead of 16.4%.

The balance of trade, already showing a deficit, further deteriorated (Figure IX) because imports grew in response to increased demand from households. Moreover, given that the majority of gas consumed in France is imported,[15] the

15. Electricity, on the other hand, is mainly produced in France. Household energy consumption ratios in ThreeME are approximately 99% for locally produced electricity, while 99% of gas is imported.

Figure VIII – **Impact of the price cap: GDP and its national components**

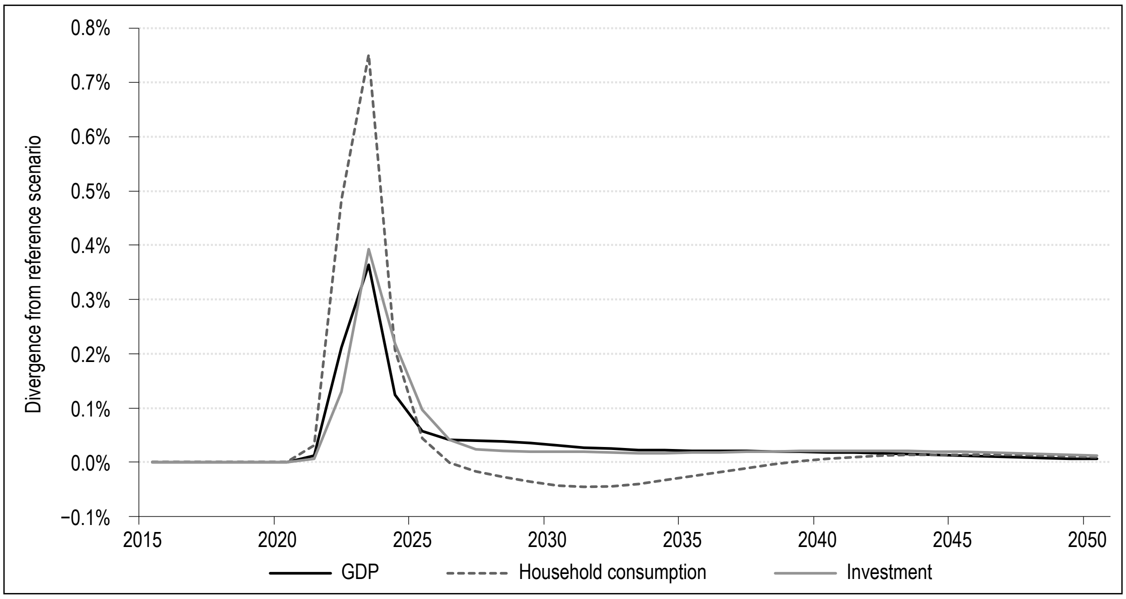

Note: The variables are given in volume terms.
Source: ThreeME simulations.

Figure IX – **Balance of foreign trade by product, expressed in percentage points of nominal GDP (absolute divergence from reference scenario)**

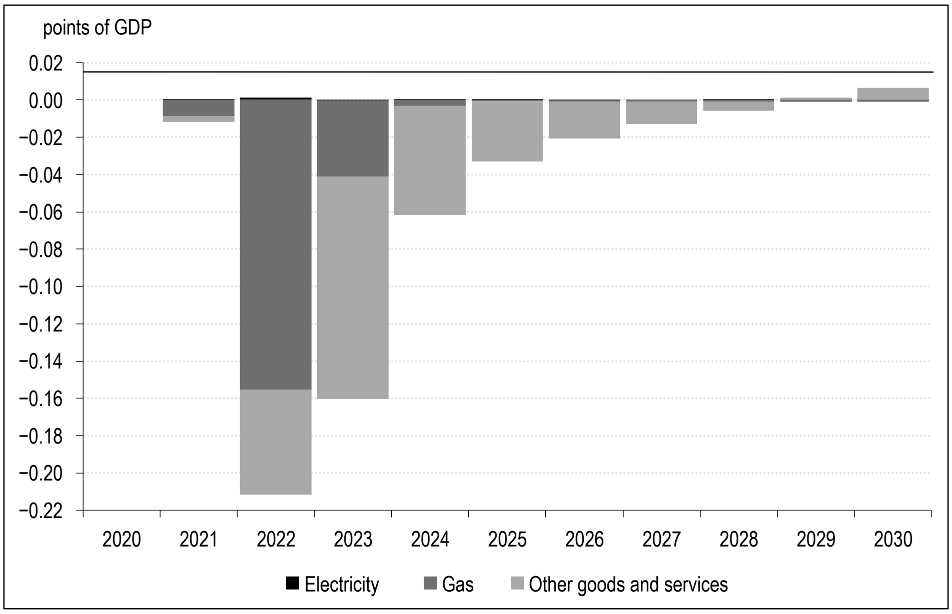

reduction in consumer prices brought about by the price cap actually led to an increase in demand. As a result, in comparison with the reference scenario, our results show that the trade deficit widened by 0.21 GDP points in 2022 and 0.16 points in 2023.

Medium-term variations in imports and exports compared with the reference scenario are detailed in Appendix 1 (see Figure A1-II). They show a peak increase in imports when the price cap is in place, then a post-crisis increase in exports thanks to more favourable terms of trade made possible by domestic price subsidies. Nonetheless, the balance of trade (Figure IX) continues to deteriorate, albeit to a lesser extent, in comparison to the reference scenario. This decline stops in 2030.[16]

3.1. Estimating the Budgetary Cost

With regard to public finances, nominal public spending increases by 0.7% in 2022 and 0.8% in 2023, largely as a result of the cost of this policy. In 2022, these simulations estimate the cost of the policy at approximately 0.6% of nominal GDP, with the cost of the electricity price cap accounting for 15% of the total cost, and the gas price cap accounting for the rest (i.e. 85%, see Figure XI).

In the following year, the cost of the electricity price cap increases significantly. This increase results from the inclusion, in the regulated price, of costs intended to offset the high prices seen on the wholesale gas market, which had not been taken into consideration for the preceding year. As a result, the total cost of the price cap policy stands at around 0.7% of nominal GDP for the year in question, with electricity accounting for approximately 90% of that cost.

For natural gas, however, falling prices on the wholesale market serve to drive down the cost of the price cap policy.[17] As noted by the CRE, even though the theoretical gas price could be even lower, it is fixed at the regulated price for Q2 2023, which has the effect of offsetting the cost of the policy for the rest of the year.[18] This decrease can also be partly attributed to the end of the gas price cap policy in June 2023, and the end of regulated pricing.

On account of the cost of the price cap policy, the public spending deficit (Figure X) deteriorates further in comparison to the reference scenario, which already incorporates the shock associated with the slowdown of economic activity.

16. In level terms, the trade balance remains in deficit. In the longer term (2050), the trade deficit in the price cap scenario becomes larger than that in the reference scenario.
17. We developed an alternative scenario in which the theoretical gas price was set based on the decline in TTF prices, while having a policy forcing the consumer price to be 15% higher than the previous year, as dictated by the cap price policy. This configuration ultimately resulted in a regulated price higher than the theoretical price, thus creating a negative cost (i.e. a revenue) for the gas policy.
18. The gas policy is only applicable in the first half of 2023, which further reduces its cost.

Figure X – **Primary deficit to GDP ratio, for each scenario**

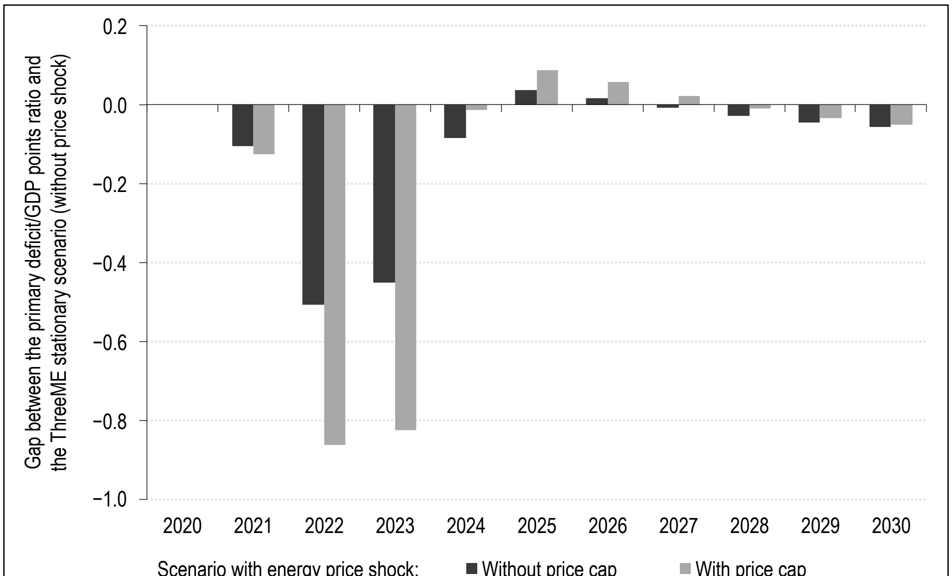

Note: The primary deficit/GDP ratio in the ThreeME stationary scenario is −0.9 point of GDP. As such, even when the difference with the latter is positive, as it is in 2025, the primary balance remains clearly in deficit.
Source: ThreeME simulations.

Figure XI – **Estimating the cost of the energy price cap policy**

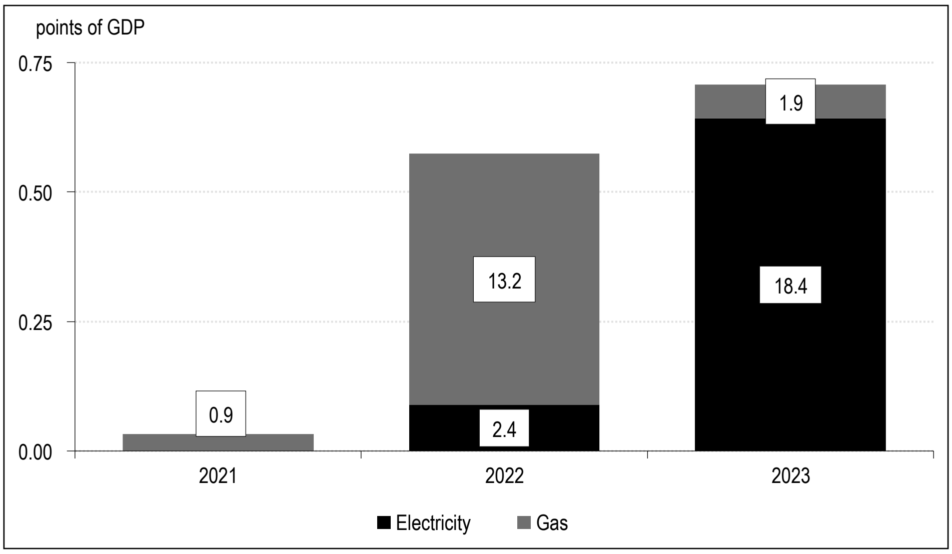

Note: The values shown in Euros were computed in the context of the ThreeME simulations and are not directly comparable with the real costs.
Source: ThreeME simulations.

In the context of the model used here, interest rates are regarded as exogenous factors which do not respond to inflation. As such, although inflation has the effect of worsening the deficit during the crisis, the fall in prices observed post-crisis has a positive impact on the accumulated debt burden when calculating the primary deficit. Nevertheless, this effect is fleeting. Subsequently, we can observe a second, slight but enduring deterioration in the deficit in terms of nominal GDP. This can be attributed to an insufficiently robust economic rebound.

For the period to 2030 the budgetary coefficient remains below one, dipping to 0.6. The reduction in inflation made possible by the price cap measure, combined with the boost to economic activity provided by the decision to shield energy prices, are not sufficient to offset the long-term cost of the policy.

During the period when the price cap was in place, we noted an overall reduction of inflation by 0.6 pp in 2022 and 0.4 pp in 2023 (Figure XII). In 2024, energy prices rose more

Figure XII – **Impact of the price cap on inflation, by contribution (divergence from reference scenario)**

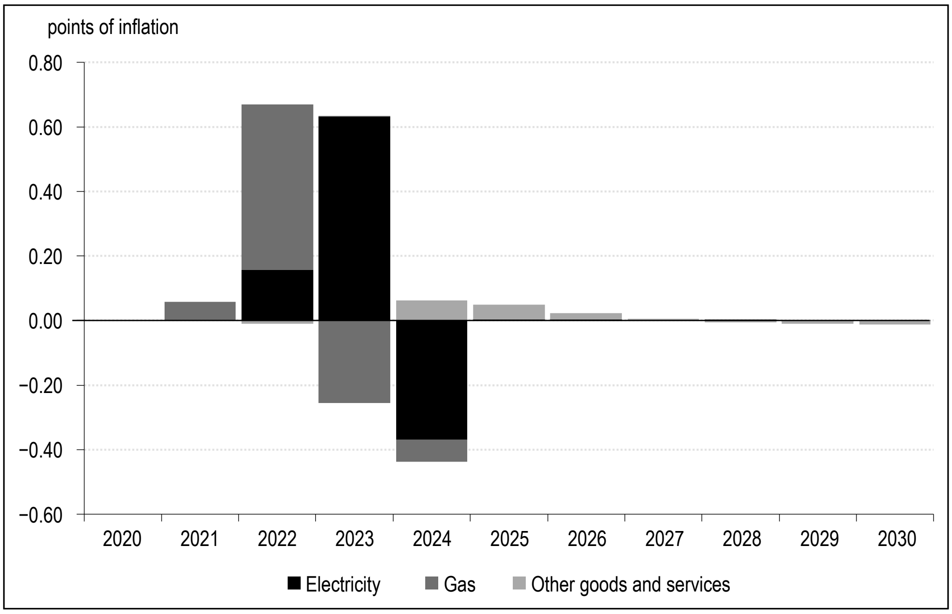

Source: ThreeME simulations.

sharply in our model than in the reference scenario, contributing to an overall 0.4 pp increase in inflation (the contribution of energy products to the inflation differential is 0.4 pp). The price cap also serves to mitigate inflation in the cost of goods and services other than energy, with inflation falling to a lesser extent during the period in question.

Since the purpose of the price cap policy was to protect household purchasing power by mitigating the rise in energy prices, the consequences for emissions are, unsurprisingly, similar to those recorded for energy consumption, namely an increase in direct greenhouse gas emissions from households in comparison with the reference scenario, peaking at 2.5% in 2023 and then falling back to their initial level (see Figure XIII).

Figure XIII – **Direct household emissions (divergence from reference scenario)**

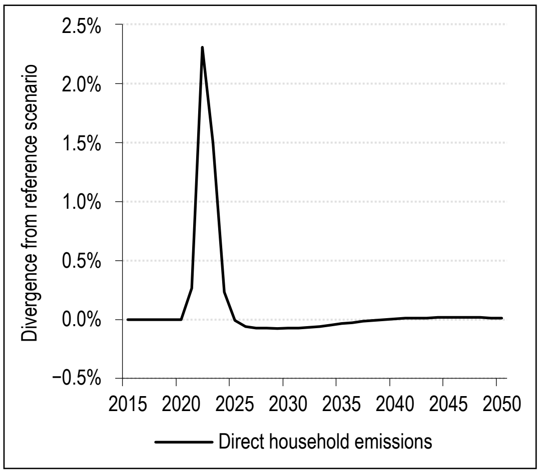

Source: ThreeME simulations.

3.2. Sensitivity Tests

The results of our simulations are immediately dependent upon the distribution of parameters used when calibrating the model, both in terms of sensitivity to price changes (price elasticity) and the speed at which such reactions occur (adjustment parameters). In order to assess the sensitivity of our macroeconomic results to the choices made regarding these parameters, we simulated multiple variants of the same price cap scenario, each with a different value for substitution elasticity η^{LESCES} in the equation for marginal propensity to consume (Equation 24 in Appendix 2). In our central scenario, the price-elasticity value is equal to –0.5, and in this section we test six alternative values {–1; –0.8; –0.6; –0.4; –0.2; 0}.

We thus see that the effect of the price cap policy on consumption sits somewhere between 0.39% (elasticity at –1) and 0.64% (elasticity at 0) for 2022, and between 0.57% and 1.03% for 2023, converging very rapidly around the central value for the effects actually observed once the policy was rescinded (Figure XIV). The effect measured in terms of GDP variation remains similar, ranging from 0.48% (price elasticity at –1) and 0.28% (price elasticity at 0). These tests tell us that although the amplitude of the

Figure XIV – **Impact on household consumption and GDP, applying multiple price elasticity values to the household demand function**

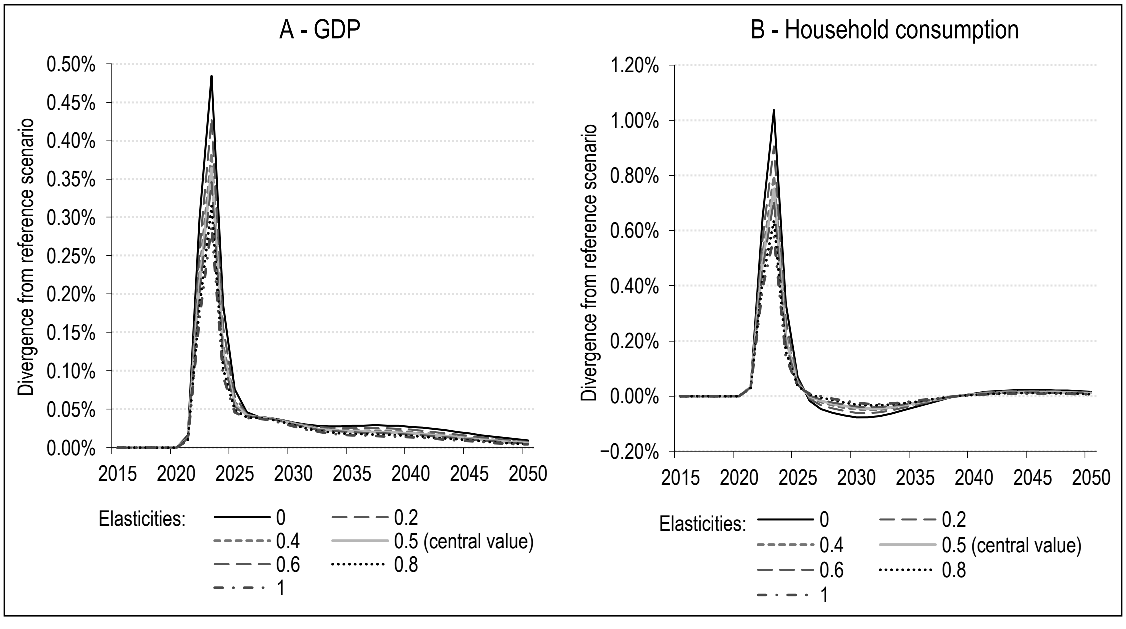

Source: ThreeME simulations.

effect observed is heavily dependent on the choice of value utilised for this parameter, the macroeconomic dynamics we identified remain valid nonetheless.

4. Discussion

The price cap policy was remarkably efficient at achieving its initial objectives, particularly protecting consumer purchasing power and keeping the high rate of inflation under control. In 2022 and 2023, a significant reduction in energy consumption was observed, which can be ascribed not only to rising prices and favourable meteorological conditions, but also to the context of uncertainty surrounding future price developments.

Regarding inflation, the National Institute of Statistics and Economic Studies (INSEE) estimates that this policy succeeded in cutting France's rate of inflation in half (Bourgeois & Lafrogne-Joussier, 2022).

In terms of all-round efficacy, policies of this nature raise several questions, with regard to both their long-term economic implications and their structural consequences. One of the great disadvantages with public policies of this kind is that they neutralise price signals, and thus do not necessarily encourage households to restrict their consumption as much as they would have done in the absence of such measures. By way of a comparison, the German system more effectively integrates the price signal dimension by protecting only a certain proportion of consumption,[19] with the rest remaining subject to market prices.[20]

France opted to implement a policy which was beneficial to all consumers, as when the government took measures to reduce fuel prices, with no means-testing of these benefits. As noted in our results, household consumption was higher when the price cap was in place. Making this a means-tested measure would avoid the risk of subsidising the consumption of households who are relatively impervious to fluctuations in energy prices, while helping low-income households to escape energy poverty, in line with the conclusions reached by Chaton & Gouraud (2020). In order to simulate such specifications in ThreeME, we would have needed to adjust the model so as to split households into categories based on their level of income. This is not possible with the version of the model we used. Similar, or even better, results would most likely have been obtained if household resources could have been taken into consideration as a condition.

We have already noted that the cost of the policy depended not only upon the pricing system, but also, as expected, on gas prices on the wholesale market. These prices are highly volatile (as demonstrated by Figure III), which means that predictions for the final cost of the policy are highly uncertain and dependent upon the current price at time of calculation. The French Finance Ministry has issued a series of forecasts at approximately six months intervals as part of their annual budget calculations, starting with the annual budget forecast for 2023, published in 2022. The most recent available figures come from the 2024 Stability Programme, published in April 2024. The decline in gas spot prices observed since late August 2022 has drastically reduced the overall cost of the price cap policy, when compared with the government's initial estimates. For the year 2023, for example, the initial estimate for the cost of the gas price cap was 9.8 billion euros; this was when Dutch TTF prices were at their peak. This estimate was revised to just 2 billion euros in April 2024, following a significant fall in prices (cf. Figure IV). The Dutch TTF returned to levels comparable to those last seen before the energy crisis.

While this may not appear to be a problem *per se* – in principle, less costly policies are always welcome, especially coming after several years of unscheduled spending increases – it is nonetheless difficult to ignore, particularly in these uncertain times, the possibility that things could go the other way, i.e. that energy prices could remain high for an extended period of time, thus exacerbating the burden placed on public finances.

This raises questions as to how political decision-makers handle unscheduled expenditure, with potential consequences including a blow to the credibility of their economic policies, as well as unexpected increases in government debt. We might also suggest that overestimating the budgetary cost of policies could have a crowding out effect on the public finances, ultimately influencing the allocation of funds between different policies within the budget, and undermining the efficiency of this allocation.

The primary aim of the research presented here is to evaluate a policy which involved capping prices during a time of crisis. Nonetheless, this evaluation runs up against certain limitations

19. 80% of past consumption.
20. However, it should be noted that there is no distinction between off-peak and peak consumption, and therefore no incentive to moderate consumption when production costs are highest.

arising from the modelling framework used. It is important to bear these limitations in mind.

When evaluating the cost of the price cap policy, we note that our results differ from the government's own estimates (see Figure IV), even when we remove the effective cost of the fiscal exemptions, which were not taken into account by ThreeME. This suggests that the government estimates might be based on different estimates for theoretical energy prices. In the most recent government estimates (PSTAB 2024), the "compensation to suppliers" paid under the gas price cap scheme is estimated at 4.5 billion euros in 2022 and 2 billion in 2023 (compared with 13.2 billion and 1.9 billion in our simulations). For the electricity price cap, the government estimates are 10.3 billion euros for 2022 and 15.5 for 2023 (compared with 2.4 billion and 184 billion in our simulations). It thus appears that our simulations overestimate the cost of the gas price cap in 2022, and yield a different picture for the distribution over time of the cost of the electricity price cap.

Nonetheless, the sum totals for these two years remain comparable, albeit slightly overestimated: for 2022 they stand at 14.8 billion euros (PSTAB 2024) compared with our 15.6 billion (ThreeME), and for 2023 the figures are 17.5 billion euros (PSTAB 2024) to our 20.3 billion (ThreeME). Our study does not include one component of the price cap scheme, namely the partial tax exemptions on gas and electricity which also helped to drive prices down. We should also note that, although the model appears to behave in a largely linear manner, it is not necessarily possible in this specific case to estimate by extrapolation the effect of the price cap policy as a whole, i.e. including the tax exemptions not taken into consideration by this study. These tax breaks applied only to the electricity price cap, while our simulations offer combined results for both gas and electricity.

Another limitation of this study is the way it models foreign trade. Our model is constructed to study France, and although the French economy is open this is only represented by a "rest of world" component, with little responsiveness apart from global demand which responds to fluctuations in relative prices. Within the context of this analytical exercise, the rest of the world is not affected by the energy price shock: the analytical framework therefore does not accurately represent all of the shocks engendered by the sudden spike in energy prices in Europe. One of the direct consequences of this increase might be a decrease in the demand for French goods from France's trading partners, first and foremost fellow EU members, which would theoretically exacerbate the trade deficit by accelerating the decline in exports. The political responses of other countries to the price shock are also overlooked, despite the fact that France's principal European partners also introduced measures to attenuate price increases and protect their economies, measures which had knock-on effects for energy prices, as demonstrated by Bayer *et al.* (2023). For these reasons, the results for foreign trade could be less positive for the French economy.

* *
*

Our simulations show that the price cap system succeeded in protecting economic activity: compared to a reference scenario without price capping, GDP was 0.2% higher in 2022 and 0.4% in 2023, at a budgetary cost estimated at 0.5% of nominal GDP in 2022 and 0.7% in 2023. We estimate the budgetary coefficient for the eight years following the introduction of this policy at 0.6. Moreover, household consumption in 2022 was 0.75% higher than it would have been without the measure.

It is, however, important to bear in mind the emergency context in which these policies were introduced and the speed with which they were rolled out, largely by replicating an existing instrument which had previously been implemented in overseas regions and *départements*.

As such, although the price cap scheme succeeded in keeping inflation at manageable levels, particularly for those households most exposed to it, it is nonetheless difficult to regard it as a viable long-term policy, as it pays too little heed to the efficacy of public spending, and seems at odds with policies aimed at speeding the transition to a low-carbon economy. Introducing an indirect subsidy for final energy consumption has the effect of scrambling the price signal, a signal which could have encouraged greater moderation in energy consumption.

This research makes use of a modelling approach (the ThreeME model) capable of representing the sectoral specificities of a targeted price shock. Our work could be expanded to fully integrate the other components of the price cap policy – such as direct transfers to households or tax exemptions on energy products – in order to more precisely calculate their budgetary cost. Evaluating public policy with the help of a model which combines a climate and energy component with a multisectoral, macroeconomic analytical framework also highlights how important it is for institutions to better integrate these considerations into the traditional mechanisms of policy-making, even when the policies in question are designed to respond to short-term challenges. ☐

BIBLIOGRAPHY

Ari, A., Arregui, N., Black, S., Celasun, O., Iakova, D., Mineshima, A., Mylonas, V., Parry, I., Teodoru, I. & Zhunussova, K. (2022). Surging Energy Prices in Europe in the Aftermath of the War: How to Support the Vulnerable and Speed up the Transition Away from Fossil Fuels. IMF, *Working Papers* 2022/152. https://www.imf.org/-/media/Files/Publications/WP/2022/English/wpiea2022152-print-pdf.ashx

Armington, P. S. (1969). A Theory of Demand for Products Distinguished by Place of Production (Une théorie de la demande de produits differencies d'apres leur origine) (Una teoria de la demanda de productos distinguiendolos segun el lugar de produccion). *Staff Papers*, 16(1), 159–178. https://doi.org/10.2307/3866403

Auclert, A., Monnery, H., Rognlie, M. & Straub, L. (2023). Managing an Energy Shock: Fiscal and Monetary Policy. NBER, *Working Papers* 31543. https://ideas.repec.org/p/nbr/nberwo/31543.html

Bayer, C., Kriwoluzky, A., Müller, G. J. & Seyrich, F. (2023). Hicks in HANK: Fiscal Responses to an Energy Shock. CRC TR 224 *Discussion Paper Series* crctr224_2023_474. University of Bonn; University of Mannheim, Germany. https://ideas.repec.org/p/bon/boncrc/crctr224_2023_474.html

Blanchard, O. & Katz, L. F. (1999). Wage Dynamics: Reconciling Theory and Evidence. *American Economic Review*, 89(2), 69–74. https://doi.org/10.1257/aer.89.2.69

Bourgeois, A. & Lafrogne-Joussier, R. (2022). La flambée des prix de l'énergie : un effet sur l'inflation réduit de moitié par le « bouclier tarifaire ». *Insee Analyses* N° 75. https://www.insee.fr/fr/statistiques/6524161

Brown, M. & Heien, D. (1972). The S-Branch Utility Tree: A Generalization of the Linear Expenditure System. *Econometrica: Journal of the Econometric Society*, 40(4), 737–747. https://doi.org/10.2307/1912967

Callonnec, G. & Cancé, R. (2022). Évaluation macroéconomique de la Stratégie nationale bas-carbone (SNBC2) avec le modèle ThreeME. Ministère de la Transition Écologique, *Document de travail*, février 2022. https://www.ecologie.gouv.fr/sites/default/files/documents/%C3%89valuation%20macro%C3%A9conomique%20de%20la%20Strat%C3%A9gie%20nationale%20bas-carbone_0.pdf

Callonnec, G., Landa, G., Malliet, P. & Reynes, F. (2013). Macro-economic assessment of energy visions 2030-2050 by the ADEME-Technical document. ADEME. https://librairie.ademe.fr/ged/6711/evaluation-macroeconomique-visions-energetiques-2030-2050-med00090136.pdf

Chaton, C. & Gouraud, A. (2020). Simulation of fuel poverty in France. *Energy Policy*, 140, 111434. https://doi.org/10.1016/j.enpol.2020.111434

Fishbone, L. G. & Abilock, H. (1981). Markal, a linear-programming model for energy systems analysis: Technical description of the BNL version. *International Journal of Energy Research,* 5(4), 353–375. https://doi.org/10.1002/er.4440050406

Hamdi-Cherif, M., Malliet, P., Plane, M., Reynes, F., Saraceno, F. & Tourbah, A. (2022). *Greening Europe: 2022 European Public Investment Outlook*. Open Book Publishers. https://www.openbookpublishers.com/books/10.11647/obp.0328

Hernnäs, H., Johannesson-Lindén, Å., Kasdorp, R. & Spooner, M. (2023). Pass-through in EU electricity and gas markets. *Quarterly Report on the Euro Area (QREA),* 22(2), 23–34. https://ideas.repec.org/a/euf/qreuro/0222-02.html

Heyer, E., Reynès, F. & Sterdyniak, H. (2007). Structural and reduced approaches of the equilibrium rate of unemployment, a comparison between France and the United States. *Economic Modelling,* 24(1), 42–65. https://doi.org/10.1016/j.econmod.2006.06.005

IEA (2023). *Europe's energy crisis: Understanding the drivers of the fall in electricity demand.* https://www.iea.org/commentaries/europe-s-energy-crisis-understanding-the-drivers-of-the-fall-in-electricity-demand

Kehoe, T. J. (1996). Social Accounting Matrices And Applied General Equilibrium Models. Federal Reserve Bank of Minneapolis, *Working Paper* 563. https://users.econ.umn.edu/~tkehoe/papers/wp563.pdf

Klein, C. & Simon, O. (2010). Le modèle MÉSANGE : nouvelle version réestimée en base 2000. *Document de Travail de la DGTPE* N° 2010/02, 66–205.
https://www.tresor.economie.gouv.fr/Articles/c8c71bdc-bc87-4430-b593-ea0ba60c01cf/files/85b1504b-4967-4f7d-92b0-ec96b4d920df

Langot, F., Malmberg, S., Tripier, F. & Hairault, J.-O. (2023). The Macroeconomic and Redistributive Effects of Shielding Consumers from Rising Energy Prices: the French Experiment. CEPREMAP, *Working Papers* 2305. https://ideas.repec.org/p/cpm/docweb/2305.html

Layard, R., Nickell, S. J. & Jackman, R. (2005). *Unemployment: Macroeconomic performance and the labour market*. Oxford University Press on Demand.

Loulou, R., Remme, U., Kanudia, A., Lehtila, A. & Goldstein, G. (2005). *Documentation for the TIMES model part II*. Energy Technology Systems Analysis Programme.

Malliet, P., Reynès, F., Landa, G., Hamdi-Cherif, M. & Saussay, A. (2020). Assessing short-term and long-term economic and environmental effects of the COVID-19 crisis in France. *Environmental and Resource Economics,* 76(4), 867–883. https://link.springer.com/article/10.1007/s10640-020-00488-z

Reynès, F. (2010). The Phillips curve as a more general model than the Wage Setting curve. OFCE, *Working Paper* 2010-28.

Reynès, F. (2019). The Cobb–Douglas function as a flexible function: A new perspective on homogeneous functions through the lens of output elasticities. *Mathematical Social Sciences,* 97, 11–17.
https://doi.org/10.1016/j.mathsocsci.2018.10.002

Reynès, F., Callonnec, G., Saussay, A., Landa, G., Malliet, P., Gueret, A., Hu, J., Hamdi-Cherif, M. & Gouëdard, H. (2021). *ThreeME Version 3 Multi-sector Macroeconomic Model for the Evaluation of Environmental and Energy Policy: A full description*. February 2021.

Sgaravatti, G., Tagliapietra, S., Trasi, C. & Zachmann, G. (2023). National policies to shield consumers from rising energy prices. *Bruegel Datasets*.
https://www.bruegel.org/dataset/national-policies-shield-consumers-rising-energy-prices

RESULTS

Figure A1-I – **Household energy consumption**

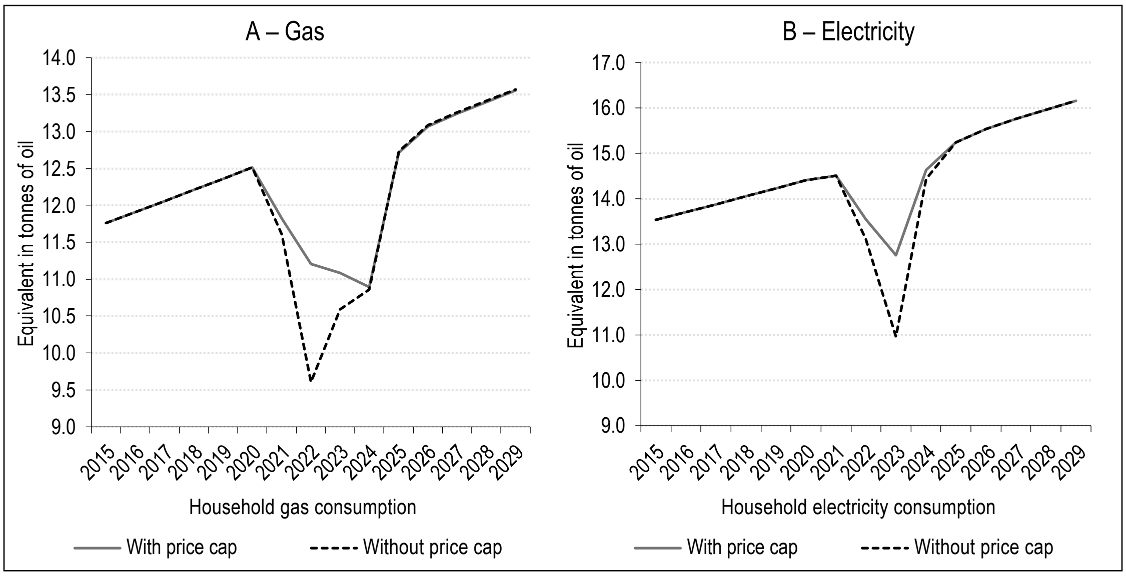

Source: ThreeME simulations.

Figure A1-II – **Impact of the price cap on foreign trade (divergence from reference scenario)**

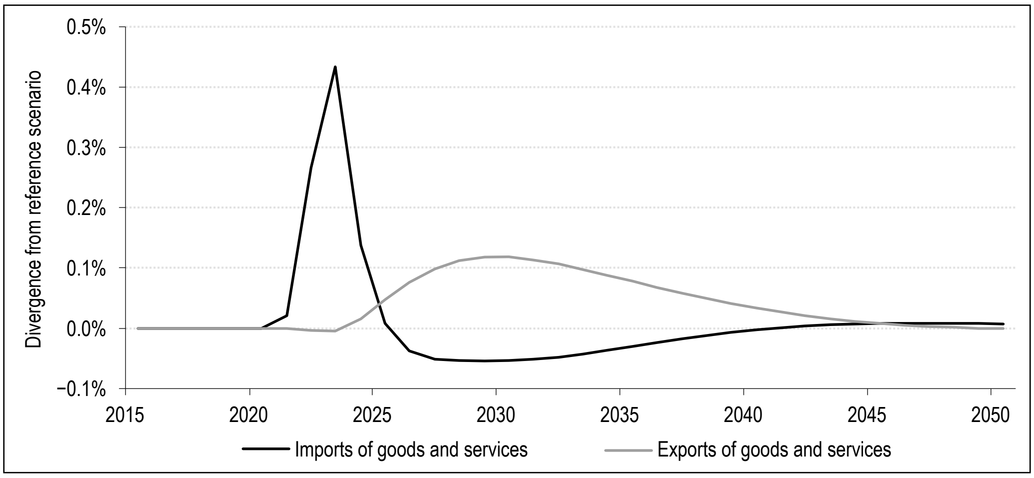

Source: ThreeME simulations.

APPENDIX 2

THE PRINCIPAL EQUATIONS USED BY ThreeME

Specifications of the Adjustment Mechanisms

Unlike Walras-inspired models which presuppose that supply and demand are always perfectly balanced, with perfect flexibility of both prices and quantities, ThreeME takes a more realistic view of the way in which the economy functions, explicitly incorporating the gradual adjustment of prices and quantities (factors of production, consumption). Within this Keynesian framework, permanent or transitory balances of under-employment are possible, with demand shaping supply. ThreeME supposes that real levels of prices and quantities gradually adjust to their notional levels. The notional level corresponds to the optimal level (desired, or targeted) that the economic agent in question (businesses for prices and demand for factors of production, households for consumption, the central bank for interest rates, etc.) would choose if adjustment constraints were not an issue. These constraints primarily arise from adjustment costs, physical or temporal limitations and other sources of uncertainty. In formal terms, we assume that the adjustment process and forward planning regarding prices and quantities can be represented by the following equations:

$$\log X_t = \lambda_0^X \log X_t^n + \left(1-\lambda_{0,t}^X\right)\left(\log X_{t-1} + \Delta\left(\log X_t^e\right)\right) \tag{2}$$

and

$$\Delta\left(\log X_t^e\right) = \lambda_t^{1,X}\,\Delta\left(\log X_{t-1}^e\right) + \lambda_t^{2,F}\,\Delta\left(\log X_{t-1}\right) + \lambda_t^{3,X}\,\Delta\left(\log X_t^n\right). \tag{3}$$

Where X_t is the real value of a given variable (for example, production price, labour, capital, etc.), X_t^n represents its notional level, X_t^e is its expected value for the period t, and α_i^X are the adjustment parameters (and $\alpha^{1,X} + \alpha^{2,X} + \alpha^{3,X} = 1$).

Equation (2) supposes a geometric process of adjustment. Taking expectations into account ensures that the real variables converge towards their desired long-term levels. Equation (3) supposes that these expectations are adaptable (retrospectively). It is worth noting that equation (2) and equation (3) can be reformulated within an error correction model, as used to produce econometric estimates, in order to account for the non-stationary nature of certain variables:

$$\Delta\log\left(X_{t-1}\right) = \alpha_1\,\Delta\log\left(X_{t-1}\right) + \alpha_2\,\Delta\log\left(X_{t-1}^n\right) - \alpha_3\log\left(X_{t-1}\right)/\left(X_{t-1}^n\right).$$

To do so, the following constraints must be respected:

$\lambda_0^X = \alpha_3, \lambda_1^X = 0, \lambda_2^X = \alpha_1/\left(1-\alpha_3\right), \lambda_3^X = \left(\alpha_2-\alpha_3\right)/\left(1-\alpha_3\right).$

We also suppose that the substitution effects $\left(SUBST_X\right)$ adjust slowly to the notional substitution effects $\left(SUBST_X^n\right)$:

$$SUBST\_X_t = \lambda_4^X\, SUBST\_X_t^n + \left(1-\lambda_4^X\right)*SUBST\_X_{t-1}. \tag{4}$$

The three equations shown above allow for a broad array of adjustments, as they integrate different forms of rigidity (on prices and quantities, on expectations and on substitution mechanisms). By way of an example, let us consider the full specification for demand for labour (L). For simplicity's sake, the sectoral index is omitted. Notional demand for labour (L^n) can be derived by minimising production costs. It is positively dependent upon production levels $\left(Y\right)$, negatively dependent upon the productivity of labour ($PROG_L$) and another component combining all of the substitution phenomena with the other factors of production ($SUBST_L$):

$$\Delta\log\left(L_t^n\right) = \Delta\log\left(Y_{t-1}\right)\Delta\log\left(PROG\_L_t\right) + \Delta SUBST\_L_t. \tag{5}$$

We introduce a distinction between the real and notional substitution effects, in order to take account of the fact that demand for labour generally responds more rapidly to changes in the level of production than to substitution phenomena: while it is physically necessary to use more labour to respond to an increase in production, substitutions imply making changes to the structure of production, and implementing these changes may take longer. Real substitution thus adjusts gradually to notional substitution $\left(SUBST_L^n\right)$, which depends on the relative prices of the factors of production:

$$\Delta SUBST\_L_t^n = -\eta^{LK}\varphi_{t-1}^K\Delta\log\left(C_t^L/C_t^K\right) - \eta^{LE}\varphi_{(t-1)}^E\Delta\log\left(C_t^L/C_t^E\right) - \eta^L\varphi_{t-1}^M\Delta\log\left(C_t^L/C_t^M\right) \tag{6}$$

where η^{LK}, η^{LE}, η^{LM} are the substitution elasticities between labour and the other factors of production, capital, energy and materials (i.e. intermediate consumption not related to energy) respectively. φ^K, φ^E, φ^M represent capital, energy and materials, respectively, proportionally to total production costs. C^K, C^L, C^E, C^M are the unit production costs of capital, labour energy and materials, respectively. In the following section, we provide more details regarding the derivation of demand for these factors. Finally, as the adjustment mechanisms are defined by means of equations (4), (5) and (6), the three following ratios are used:

$$\log\left(L_t\right) = \lambda_0^L\log\left(L_t^n\right) + \left(1-\lambda_0^L\right)\left(\log\left(L_{t-1}\right) + \Delta\log\left(L_t^e\right)\right)$$

$$\Delta\log\left(L_t^e\right) = \lambda_1^L\,\Delta\log\left(L_{t-1}^e\right) + \lambda_2^L\,\Delta\log L_{t-1}) + \lambda_3^L\,\Delta\log\left(L_t^n\right) \tag{7}$$

$$SUBST_{L_t} = \lambda_4^L\, SUBST_{L_t^n} + \left(1-\lambda_4^L\right)SUBST_{L_{t-1}}. \tag{8}$$

Production Functions and Demand for Factors of Production

The structure of production is broken down into three levels (see Figure A2). The first supposes a function of production operating with four inputs (or factors of production), often denoted by the acronym KLEM (capital, labour, energy and materials). The first level also incorporates a fifth element: transport and commercial margins. Strictly speaking, the latter should not be considered as factors of production, since they come into play after the production process. As such, they cannot be substituted for the factors of production. Nevertheless, they are inextricably linked with production levels, since once goods have been manufactured they need to be transported and brought to market. At the second level, aggregates for investment, energy, materials and margins are broken down with reference to the type of products involved (e.g. energy sources). At the third level, demand for each factor or margin is either imported or produced locally. Demand for factors of production is derived by minimising companies' production costs. We assume that the function of production has constant returns to scale, more general than CES (constant elasticity of substitution), to the extent that substitution elasticities may vary from one pair of inputs to the next (Reynès, 2019). Minimising production costs gives us the following equations for notional demand for factors of production. This is applicable to any and all economic activities, but in order to simplify the algebra the sectoral index is omitted here:

$$\Delta\log\left(FP_{j,t}^{n}\right)=\Delta\log\left(Y_{t}\right)-\Delta\log\left(PROG\_FP_{j,t}\right)+\Delta SUBST\_FP_{j,t} \tag{9}$$

$$\Delta SUBST\_FP_{j,t}^{n}=-\sum_{\substack{j'=1\\ j\neq j'}}\eta_{j,j'}\,\varphi_{t-1}^{j'}\,\Delta\log\left(\frac{C_{(j',t)}^{FP}}{C_{j,t}^{FP}}\right), \tag{10}$$

where $\varphi_{j,t-1}=\left(C_{j,t}^{FP};FP_{j,t-1}\right)/\left(\sum_{j}C_{j,t}^{FP};FP_{j,t-1}\right)$ and $j=K,L,E,M$

where FP_{j}^{n} is the notional demand for an input j, $\eta_{j,j'}$ is the substitution elasticity between pairs of inputs j and j', $PROG\_FP_{j,t}$ is technical progress relevant to this input j, C_{j}^{FP} is the cost/price of the input j and Y is the level of production for the sector in question.

Figure A2 – **Structure of production in the ThreeME model**

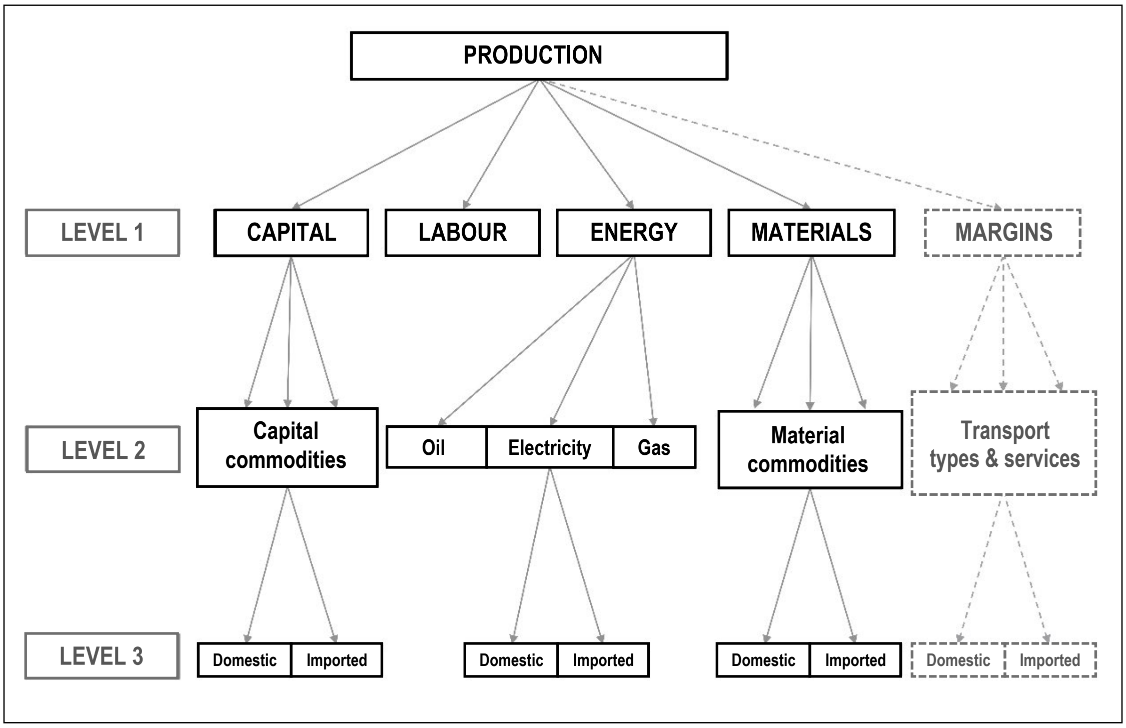

In keeping with the data from the national accounts, ThreeME supposes that all commodities could be produced by more than one sector. For example, electricity can be produced by several sectors: nuclear, wind power, etc. The output from each sector is defined by the following equations:

$$Y_{c,a}=\varphi_{c,a}YQ_{c} \tag{11}$$

$$Y_{a}=\sum_{c}Y_{a,c} \tag{12}$$

where YQ_{c} is aggregated domestic production of the commodity c. It is determined by demand (final and intermediate consumption, investment, public spending, export and inventory variation). $\varphi_{c,a}$ Is thus the proportion of the commodity c produced by the sector a (where $\sum_{a}\varphi_{c,a}=1$) and Y_{a} is the aggregated output of the sector a.

Capital and Investment Equations

In ThreeME, investment depends on expected production, past trends, substitution phenomena and a corrective mechanism, which ensures that companies reach their notional level of fixed capital stock in the long term. Fixed capital stock is subtracted from investment using the standard capital accumulation equation.

$$\Delta\log\left(IA_t\right)=\theta_1^{IA}\,\Delta\log\left(Y_t^e\right)+\theta_2^{IA}\,\Delta\log\left(IA_{t-1}\right)+\theta_3^{IA}\times d\left(SUBST\_K\right)+\theta_4^{IA}\left(\log\left(K_{t-1}^n\right)-\log\left(K_{t-1}\right)\right) \quad (13)$$

$$K_t=\left(1-\delta^K\right)K_{t-1}+IA_t, \quad (14)$$

where IA is investment, Y^e expected production, K and K^n real and notional capital stocks, $SUBST\_K$ a variable combining the substitution phenomena between capital and the other inputs, and δ^K the rate of capital depreciation. We also impost the constraint $\theta_1^{IA}+\theta_2^{IA}=1$ in order to ensure that the stationary equilibrium path does in fact exist. This specification is a compromise between the empirically observed short-term trend and the cohesiveness of the model in the long term. As seen in the MESANGE econometric model (Klein & Simon, 2010), it is common practice to estimate an investment equation rather than a capital stock equation. There are several reasons for this. Firstly, the chronological data series for capital stock are often unreliable. Secondly, estimates often do a better job of representing short-term trends in investment. In particular, they allow us to avoid capital destruction phenomena (negative investment) which are rare in practice, as companies generally prefer to wait and allow for the technical depreciation of their equipment assets. Contrary to MESANGE, however, we suppose that investment depends on the difference between the real and notional capital stocks. This ensures that the real capital stock converges towards its notional level with time. In the long term, the model is thus consistent with the "function of production" theory which holds that there is a direct relationship between production levels and capital stock (not capital flow).

Salary Equation

Various studies have shown that theoretical arguments and empirical estimates do not make it easy to choose between two specifications. Nonetheless, specification differences have significant implications for the definition of the Non-Accelerating Inflation Rate of Unemployment (NAIRU), and thus on the inflationary tendencies and long-term properties of macroeconomic models (Blanchard & Katz, 1999). In ThreeME, we chose a general specification which incorporates the Phillips curves and *Wage Settings* (WS). This supposes that notional nominal wages $\left(W_t^n\right)$ are positively dependent upon expected consumer prices $\left(P_t^e\right)$ and the productivity of labour ($PROG\_L_t$), and negatively dependent upon the unemployment rate $\left(U_t\right)$:

$$\Delta\log\left(W_t^n\right)=\rho_1^W+\rho_2^W\,\Delta\log\left(P_t^e\right)+\rho_3^W\,\Delta\log\left(PROG\_L_t\right)-\rho_4^W U_t-\rho_5^W\Delta U_t. \quad (15)$$

Alternatively, this ratio could be identical to the Phillips curve or the WS curve, depending on the values of the selected parameters (Heyer *et al.*, 2007; Reynès, 2010). The Phillips curve applies when $\rho_4^W>0$, while the WS supposes that $\rho_4^W=0$. In order for the model to have a coherent stationary state in the long term, the WS curve must also impose the constraints identified by Layard *et al*. (2005): unit-indexing of wages to prices and productivity: $\left(\rho_2^W=\rho_3^W=1\right)$ and $\rho_1^W=0$.

Household Consumption Equation

In the standard version of the model, consumption decisions are modelled using the utility function of the *Linear Expenditures System* (LES), generalised to give non-unitary substitution elasticity between goods (Brown & Heien, 1972). Household spending on each type of good varies (more or less) proportionally to their income:

$$\Delta\beta_{c,t}^{EXP}=\left(1-\eta^{LES\_CES}\right)\Delta\frac{PEXP_{c,t}}{PEXP_t^{CES}} \quad (16)$$

$$PEXP_t^{CES}=\left(\sum_c\beta_{c,0}^{EXP}\,PEXP_{c,t}^{\left(1-\eta^{LES\_CES}\right)}\right)^{\frac{1}{1-\eta^{LES\_CES}}} \quad (17)$$

Price and Margin Rate Equations

Production prices for each sector are set at their lowest level, applying a margin to the unit cost of production (including the cost of labour, capital, energy and other forms of intermediate consumption):

$$PY_t^n=CU_t\left(1+TM_t\right) \quad (18)$$

$$\Delta\log\left(1+TM_t^n\right)=\sigma^{TM}\left(\Delta\log\left(Y_t\right)-\Delta\log\left(Y_{t-1}\right)\right) \quad (19)$$

$$TM_t=\lambda^{TM}\,TM_t^n+\left(1-\lambda^{TM}\right)TM_{t-1}, \quad (20)$$

where PY_t^n is the notional price, CU_t the unit cost of production and Y_t the level of production. TM_t and TM_t^n are the real and notional margin rates, respectively.

The notional price equation is the only price equation derived from economic behaviour: supposing that the demand upon companies is negatively correlated to their prices, we can easily demonstrate that the optimal price corresponds to a margin rate on the margin cost of production. The margin rate equation reflects the fact that returns to scale are diminishing in the short term. As a result, an unexpected increase in production leads to higher marginal production costs, and thus to higher notional prices.

Other prices are defined in accounting terms, starting with production prices and applying an adjustment process:

$$\log PY_t = \lambda_0^{PY} \log PY_t^n + \left(1-\lambda_0^{PY}\right) \log PY_{t-1} + d\log PY_t^e \tag{21}$$

$$d\log PY_t^e = \lambda_1^{PY}\ PY_{t-1}^e + \lambda_2^{PY}\ PY_{t-1} + \lambda_3^{PY}\ PY_{t-1}^n. \tag{22}$$

Household Demand Equations

In the standard version of the model, consumption decisions involving choices between different products are modelled using the utility function of the *Linear Expenditures System* (LES), generalised to give non-unitary substitution elasticity between commodities. An LES specification supposes that a certain portion of consumption (NCH_c) in the reference year is autonomous, or essential, and thus that the relationship between income and consumption is not linear. This specification makes it possible to distinguish between consumption of essential commodities and other goods and services:

$$\left(CH_c^n - NCH_c\right) PCH_c = \varphi_c^{MCH} \left(CH^{VAL} - PNCH\ NCH\right), \tag{23}$$

where $\sum \varphi_c^{MCH} = 1$ and CH_c^n corresponds to the notional level of consumption of a given commodity c, PCH_c is its price, and φ_c^{MCH} the share of non-essential consumption as a proportion of total non-essential consumption (in value terms). This ratio is constant if the substitution elasticity between commodities is equal to one (the Cobb-Douglas hypothesis). In this case, (Cobb-Douglas utility function with no autonomous spending), expenditure will fluctuate proportionally to income. If we use a CES function where the substitution elasticity is η^{LESCES}, the marginal propensity to spend will vary in response to relative prices, following the specification:

$$\Delta \log \varphi_{c,t}^{MCH} = \left(1-\eta^{LESCES}\right) \Delta\left(\log \frac{PCH_{c,t}}{PCH_t^{CES}}\right) \tag{24}$$

$$PCH^{CES} = \left(\sum_c \varphi_{c,t_0}^{MCH}\ PCH_c^{1-\eta^{LESCES}}\right)^{\frac{1}{1-\eta^{LESCES}}}. \tag{25}$$

International Trade Equations

The price of a locally-produced commodity is a weighted average of the production prices (indexed against a) which went into producing that commodity. For example, the electricity price is a weighted average of the prices charged by the various electricity producing sectors. The price paid by the end user (consumers, government, sectors, rest of the world) also includes commercial margins and transport, as well as all taxes less subsidies. Combining these prices with import prices, we get the average price for each commodity, as paid by the end user.

$$\Delta \log\left(X_{c,t}\right) = \Delta \log\left(WD_{c,t}\right) + \Delta SUBST\_X_{c,t} \tag{26}$$

$$\Delta SUBST_{X_{c,t}^n} = -\eta^X\ \Delta \log\left(\frac{P_{c,t}^X\ TC_t}{P_{c,t}^W}\right), \tag{27}$$

where $WD_{c,t}$ is global demand and $P_{c,t}^W$ the price. $P_{c,t}^X$ is the export price, which depends on production costs and reflects the price-competitiveness of domestically-produced goods. TC_t is the exchange rate; η^X is price-elasticity (presumed to be constant). We assume that substitution between domestic and imported goods is not perfect (Armington, 1969).

Demand for imported goods can be written:

$$A_{c,t}^M = \varphi_c^{A^M}\ A_c \tag{28}$$

$$\varphi_c^{A^M} = \frac{1}{1 + \frac{AD_c}{AM_{c,t_0}}\ e^{SUBST_c^{AM}}} \tag{29}$$

$$\Delta\left(SUBST_c^{AM}\right) = -\eta_c^{AM}\ \Delta\left(\log PAD_c - PAM_c\right) \tag{30}$$

and as a result:

$$A_{c,t}^D = \left(1-\varphi_c^{A^M}\right) A_c, \tag{31}$$

where $A_{c,t}$ represents the demand for each type of use (intermediate consumption, investment, consumption, public spending, exports, etc.) for a commodity c, while $P_{c,t}^A$ is the price. $A_{c,t}^M$ and $A_{c,t}^D$ are, respectively, the imported products and domestic products wanted for each type of usage A, with $P_{c,t}^{AM}$ and $P_{c,t}^{AD}$ their respective prices. The substitution elasticity η_c^{AM} for a type of use A of a given commodity c can vary, which allows for a high degree of flexibility. With regard to demand for use in intermediate consumption or investment, equations are constructed for each sector a, and these equations are specified for each $A_{c,a,t}$.

A full description of the model is available online at www.threeme.org.

Calibration of the Parameters

For the purposes of this article, the parameters were calibrated using the following values and, with the exception of the values pertaining to autonomous consumption of energy products (specific to this study), the same values were also used in Malliet *et al.* (2020).

Table A2-1 – **Calibration of behavioural parameters**

Elasticity parameters	Value
Elasticity of substitution between factors of production ($\eta_{F,F'}$)	0.5
Elasticity of substitution between energy sources ($\eta^{NRJ}_{c,c'}$)	0.2
Elasticity of substitution between modes of transport ($\eta^{TRSP}_{c,c'}$)	0.2
Elasticity of substitution between consumer goods (η^{LESCES})	0.5
Armington elasticity ($\eta_{c,c'}$)	0.8
Elasticity between margin rate and demand (σ^{TM})	0.75

Table A2-2 – **Calibration of adjustment parameters**

Adjustment parameters	Value
Price equations	
λ_0^{PY}	0.5
λ_1^{PY}	0.7
λ_2^{PY}	0.1
λ_3^{PY}	0.2
Salary equations	
ρ_1^{W}	0
ρ_2^{W}	1
ρ_3^{W}	1
ρ_4^{W}	0
ρ_5^{W}	0.6
Equations for factors of production	
λ_0^{L}	0.5
λ_0^{E}	0.9
λ_0^{M}	0.9
λ_1	0.7
λ_2	0.1
λ_3	0.3
Investment equations	
θ_2^{I}	1
θ_3^{IA}	0.5
θ_4^{IA}	0.05
Production equations	
λ_0^{Ye}	0.7
λ^{TM}	0.5
Household consumption equations	
λ_0^{CH}	0.6
λ_1^{CH}	0.7
λ_2^{CH}	0.1
λ_3^{CH}	0.2

Table A2-3 – **Calibration of autonomous energy consumption**

Proportion of autonomous consumption	Value
Electricity consumption (φ_{cele}^{NCH})	0.25
Gas consumption (φ_{cgas}^{NCH})	0.4

APPENDIX 3

AVERAGE RRP FOR GAS

Figure A3-I – **Comparing average RRP for electricity, with and without tax**

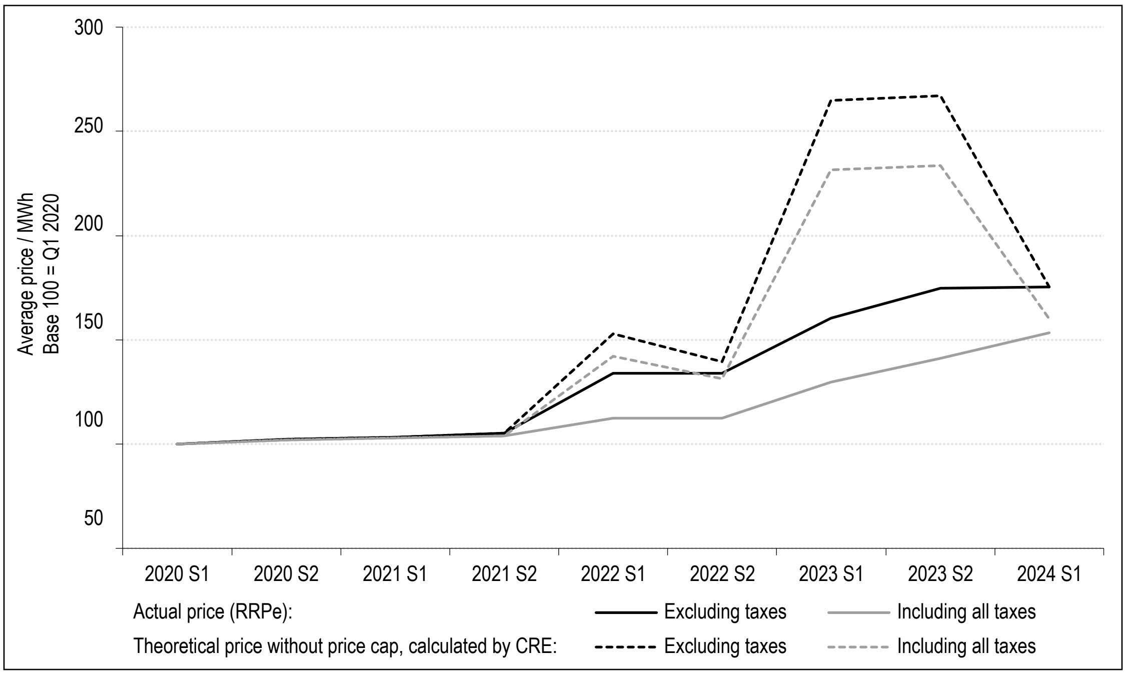

Source: CRE, authors' calculations.

Figure A3-II – **Average RRP for gas**

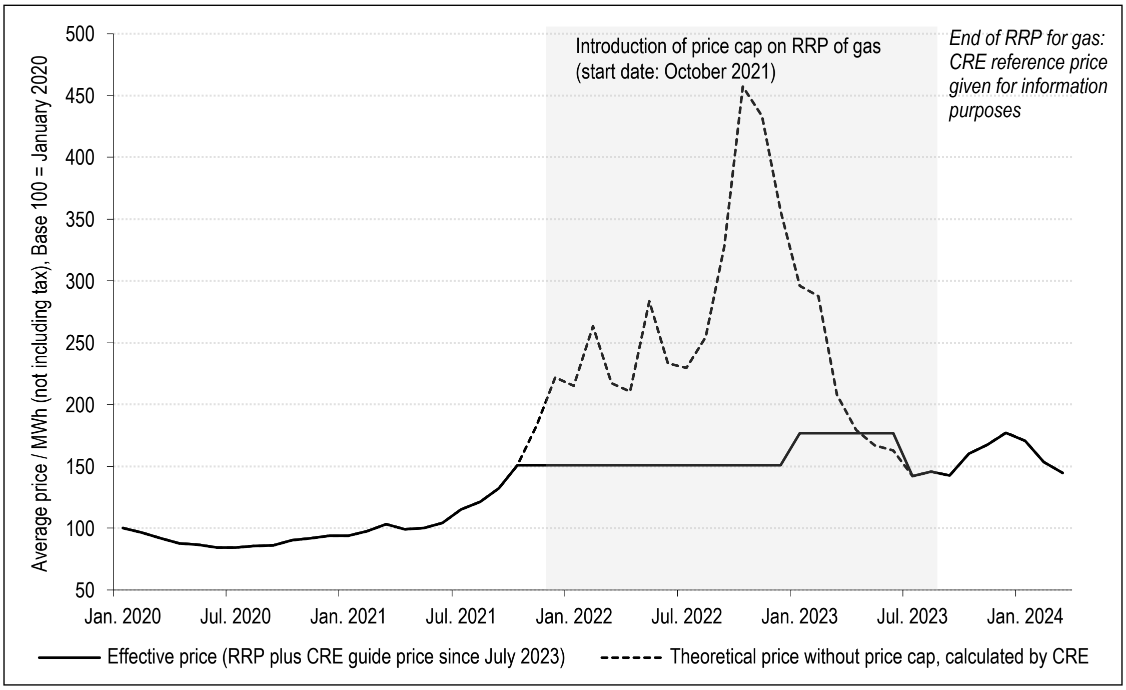

Source: CRE, authors' calculations.

Skill Distance Between Occupations and Post-Training Professional Transitions of Jobseekers

Kevin Michael Frick*, Yagan Hazard, Damien Mayaux*** and Thomas Zuber******

Abstract – Does vocational training help correct structural imbalances in the labour market? We propose a new measure of the skills distance between occupations, obtained by fine-tuning a large language model on a sample of job offers. Using this method, we demonstrate that the “return to employment” differential between jobseekers with and without training is driven by a reallocation of workers towards occupations that are very different from their previous posts in terms of the skills required. From a purely reallocative perspective, however, the return to employment differential associated with vocational training does not appear to be driven by more jobseekers moving to occupations where employers are struggling to recruit.

JEL: J62, J68, J24
Keywords: structural imbalance, training, skills

**Toulouse School of Economics; **Collegio Carlo Alberto and ESOMAS; ***Paris School of Economics; ****Banque de France.*
Correspondence: *thomas.zuber@banque-france.fr*

The opinions expressed in this study do not represent the official positions of the Banque de France, nor of Eurosystem. This research has benefited from the financial support of DARES, as part of the “Training jobseekers – Wave 3” call for research projects. The project also received a government grant via France’s National Research Agency (ANR), as part of the Future Investments programme ANR-17-EURE-001; we were also able to make use of the computing capacities of IDRIS thanks to resource allocation scheme 2022-AD011013484R1 managed by GENCI. We would like to thank Cyril Verluise for his insightful advice in the early days of this project, and Pauline Givord and Anne Bucher for their numerous comments and their help with data access, as well as all of those who contributed to the various project steering meetings.

Received in March 2024, accepted in September 2025. Translated from “Distance entre métiers et transitions professionnelles des demandeurs d’emploi après une formation”.
The opinions and analyses presented in this article are those of the author(s) and do not necessarily reflect their institutions’ or INSEE’s views.

Citation: Frick, K. M., Hazard, Y., Mayaux, D. & Zuber, T. (2025). Skill Distance Between Occupations and Post-Training Professional Transitions of Jobseekers. *Economie et Statistique / Economics and Statistics*, 547, 49–67. doi: 10.24187/ecostat.2025.547.2138

Structural imbalances between supply and demand on the labour market have the effect of increasing both the level and the stubborn persistence of unemployment. Among these imbalances, one frequently discussed hypothesis is that of a *mismatch* between workers' existing skills and the skills sought by businesses. This mismatch leads to the coexistence of residual pockets of unemployment and so-called "bottleneck" markets where businesses looking to recruit find themselves frustrated by a scarcity of qualified labour. In theory, these imbalances could be resolved by redirecting the supply of available labour towards those areas of the market experiencing recruitment difficulties, but in practice numerous factors impede such professional mobility.

Vocational training programmes targeting jobseekers aim to address these imbalances. They do so not only by facilitating access to training for jobseekers, but also by (as much as possible) steering existing training resources towards bottleneck occupations, that is, those with the most acute recruitment difficulties. The combined effect of these two aspects should in theory maximise the impact of vocational training in terms of getting people back to work. In France, the skills investment strategy (*Plan d'investissement dans les compétences*, PIC) launched in 2018 was designed with this goal in mind.

The direct effects of training on return to employment are better understood than the indirect effects of training on the balance of the labour market through occupational transitions. Numerous studies have documented the impact of training on return to employment (see for example the meta-analysis proposed by Card *et al.*, 2018). However, with a few exceptions, little academic attention has been devoted to the indirect effects that vocational training can have on the labour market by helping to reduce structural imbalances between the supply and demand of labour in different occupations (Şahin *et al.*, 2014; Barnichon & Figura, 2015; Marinescu & Rathelot, 2018). It remains particularly difficult to quantify the extent to which a given factor, such as geographic mobility, or the skills gap between occupations, can be considered responsible for the imbalances between occupations. Nevertheless, the stakes are high: in France in 2021, these structural imbalances were estimated to be responsible for around 15% of recorded unemployment (Fontaine & Rathelot, 2022).

This article investigates whether, and if so to what extent, jobseekers who have undertaken vocational training go on to make professional transitions *(i)* involving a more significant shift in skills, compared with those who have not retrained, and *(ii)* towards sectors of the labour market where the labour demand is stronger than in sectors in which they would have been able to find employment without further training. To do so, we construct a measure of the distance between occupations in terms of their skill requirements, using the text of job offers published by Pôle Emploi (the French public employment service). Our proposed methodology involves training a neural network on so-called "pretext" tasks so that the model captures content specifically relevant to skills in the corpus of job offers, mapping each offer and each occupation to a low-dimensional skills space and defining the skill distance between any two occupations as the angular distance between the corresponding points within this space.[1] The skills which constitute the dimensions of the space cannot be interpreted individually, but by construction it is possible to compare the skills vector for any two given occupations. By measuring the skill distance between occupations, we can distinguish between professional transitions between occupations demanding relatively similar skills from professional transitions to occupations whose required skills are different from those required by their previous posts. Dawson *et al.* (2021) utilise a similar methodology.[2]

The matching of administrative data provided by the DARES ForCE programme (Training, Unemployment and Employment, French *Formation, Chômage et Emploi*) allows us to track the trajectories of jobseekers who undertook vocational training between 2018 and 2020. Our study focuses on a sub-sample of jobseekers who held a stable job within the year preceding their training programme, thus satisfying the definitions of original occupation and professional transition. Comparing the professional trajectories of trained and untrained jobseekers, leveraging the large number of control variables found in administrative data, we analyse the relationship between training programmes and return to employment, the skill distance between occupations covered during the professional transitions and the resulting differentials in labour market tightness between the occupations involved in post-training professional transitions. We conducted this comparison using double/debiased machine

1. A pretext task is a task used not because it is pertinent to the final training objective for the algorithm, but because it forces the algorithm to acquire certain desirable characteristics.
2. See also Bana et al. *(2019) and Gentzkow* et al. *(2019).*

learning (Chernozhukov *et al.*, 2018) to correct for observable differences between the test group and the control group.[3] The result of this comparison can only be interpreted as a causal effect of vocational training under a conditional independence assumption (CIA) under which selection into vocational training is only driven by observable variables. This is unlikely to be true in practice. In spite of this pitfall, which is amply discussed in the literature on the evaluation of training policies, our results allow us to offer some initial insight into the under-studied connection between vocational training and efforts to address structural imbalances in the labour market.

We study the impact of training on the professional transitions of jobseekers, and certain characteristics of the jobs they find. We can thus show that the relationship between training and the return to employment almost always involves people taking up posts which are very different from their previous occupations. About the reallocation of employment, jobseekers who have retrained are more likely to undertake a professional transition, but it appears that they do not systematically move towards occupations where recruitment difficulties are more acute than they were in their original occupations. This suggests that rethinking the range of vocational training available to jobseekers, to focus more explicitly on those occupations where employers are struggling to recruit enough workers, would serve to improve the results of such training programmes in terms of job prospects. Our methodology requires us to focus exclusively on jobseekers who have held a stable job within the preceding twelve months. This sub-population is younger, more qualified and more likely to find employment than the average jobseeker registered with Pôle Emploi. In addition to the doubts over the validity of our conditional independence assumption, this necessary limitation to our sample also limits the interpretative scope of our results.

This article also contributes to the abundant literature showing that transitioning between occupations comes at a high cost in terms of human capital (Becker, 1964). Shaw (1984) was the first to illustrate the specific importance of career trajectories involving transition between different occupations, rather than different sectors, in terms of determining workers' earning potential. In the context of the German education system, Eckardt (2022) has shown that people working in fields for which they have not been specifically trained earn less money than colleagues with the appropriate qualifications. The cost in terms of lost earnings is positively correlated with the distance between occupations in terms of their skills requirements. The costs associated with forced professional transitions have also been studied, with efforts made to assess the effects of international competition on local labour markets. Traiberman (2019) for Denmark and Basco *et al.* (2025) for France have shown that when people change occupations as a result of strong international competition, the further the new occupation is from their previous occupation, the more money they lose. Moreover, Hyman (2018) has demonstrated that training policies can help workers to make professional transitions towards occupations less vulnerable to international competition. Finally, building upon Shaw's (1984) original intuition, there is now an abundant literature demonstrating that taking the multidimensional nature of skills into account (beyond a simple linear index) is crucial to understand transitions in the labour market, the matching process between candidates and employers, and wage determination (Gathmann & Schönberg, 2010; Lindenlaub & Postel-Vinay, 2021; Guvenen *et al.*, 2020; Baley *et al.*, 2022).

The rest of this article is structured as follows. We begin by introducing a new method for measuring the skill distance between occupations, based on textual data derived from job offers posted by Pôle Emploi, and showing its quantitative and qualitative validity of this method. We then present the ForCE database, and the methodological choices made when selecting the sample and variables of interest for our study. Finally, we present the results of a comparison between the career trajectories of jobseekers who have received training and those registered with Pôle Emploi who have undergone no training.

1. A New Measure of the Skill Distance Between Occupations

Our analysis of the relationship between training and the return to employment is founded upon a new measure of the distance between occupations, reflecting the differences in the skills demanded by different occupations, which can represent an obstacle to professional transitions.

There are a number of existing sources providing quality data pertaining to both occupations and professional skills. In the United States, the O*NET system provides a detailed inventory of

3. Our checks revealed that the propensity score matching method (Rosenbaum & Rubin, 1983) yields wholly comparable results.

skills and, for each occupation in the American classification system (the Standard Occupational Classification), an indicator for frequency of use and level of expected competence in appropriate skills. It also contains a table of related occupations (the Related Occupation Matrix) which gives, for each occupation, a ranked list of the top ten occupations to which workers could most easily transfer, in light of the similarity in the skills required. In France, the *Répertoire Opérationnel des Occupations et des Emplois* (Operational Register of Occupations and Jobs - ROME V3) occupies a comparable role: as well as defining 532 occupations and 14 sectors of activity, it contains a skills register which can be matched with the classification of occupations; for any given occupation, ROME can propose a list of possible occupations to which workers might be able to switch (the Mobilities function). The ROME register also allows us to broach the question of the skill gap between specific occupations, either by studying existing overlaps between the skills associated with each occupation, or else by looking at all of the professional trajectories implied by transitions to related occupations as per ROME's Mobilities function.

Nevertheless, using ROME V3 to study the impact of vocational training is subject to certain limitations. Chief among them is the fact that the professional transitions suggested in the database are not intended to be exhaustive: they omit to mention any number of transitions which might well be possible and pertinent, and say nothing at all about more difficult transitions. Furthermore, it splits suggested professional mobility opportunities into just two categories, whereas it would be more useful to have a sliding scale for the level of difficulty involved in moving from one occupation to another, particularly occupations which are relatively far removed from individuals' previous occupations, in terms of the skills required. Finally, the skills identified in ROME are often specific to just one or two occupations, which makes it impossible to distinguish between pairs of occupations which are slightly different and pairs of occupations which are completely different. These three shortcomings represent obstacles to quantitative analysis of the impact of training on professional mobility, since we start from the hypothesis that training would allow individuals to make career changes involving larger skill gaps.

In this section, we detail the development of a measure of skill distance between occupations, training a neural network on the text of job offers. The availability of a large corpus of textual data enabled us to produce a continuous, informative measure of this distance, even for pairs of occupations which are very different from one another. After describing these data, we proceed to detail the methodology used to construct our measure and the various validation exercises involved.

1.1. Data

The principal dataset used in this study is a body of text comprising more than 4 million job offers posted on the Pôle Emploi website between December 2018 and October 2020, along with the accompanying ROME codes. We used the text of these offers to pre-train our language model, then to train our neural network to predict an occupation code based on the text of a job offer. We then rebalanced the number of offers for each occupation with reference to the initial sample, so as not to introduce any training bias. We report descriptive statistics for the job offer dataset in Online Appendix S1 (link at the end of the article).

One of our objectives was to propose an alternative to existing measures based on the ROME V3 database. We thus used two fields from ROME to calculate alternative measures of skill distance between occupations, in order to test our own method.

- The "Mobilities" field assigns a numerical code to each possible pair of occupations: 1 if inter-occupation mobility seems possible without training, and 2 if inter-occupation mobility seems possible with minimal training to activate underlying skills. We used this field to produce an alternative measure of the distance between occupations, which we call *Graph distance*, i.e. the distance plotted on the oriented graph where the nodes correspond to the ROME codes, where the presence of an edge indicates that mobility is suggested in the Mobilities field, and the weighting corresponding to level 1 or 2 is defined in advance. We also used the Mobilities field when constructing our measure of the distance between occupations.

- The "Skills" structured field assigns a list of skills, both general and specific, to each of the 532 occupations in the ROME V3 classification. We used this field to produce an alternative measure of the distance between occupations, which we call the *Structured-field distance*, which is the cosine of the angle between the vectors representing the occupations in this space, where the component i is equal to 1 if the structured-field skill i is associated with this occupation, and 0 if not. This is analogous to our measure, with the difference that it is based on a representation in a space whose dimensions correspond to the total number of skills listed in the structured fields in ROME, established primarily on the basis of expert input.

In addition to these fields derived from the ROME classification, we incorporated information regarding the French labour market into the model's training dataset. To ensure that our measure of the distance between occupations only refers to the distance in terms of required skills – excluding other factors liable to influence the behaviour of agents in the labour market, such as gender stereotypes associated with different occupations, the social prestige attached to certain professions, or even differences in recruitment conditions between different occupations – we voluntarily restricted ourselves to data which we felt lent themselves to clear interpretation in terms of professional skills, especially the average level of education of workers in a given occupation,[4] as well as lists of pairs of initial and new occupations where we observed frequent transitions within companies, accompanied by wage increases, which we interpreted as instances of vertical professional mobility.[5]

1.2. The Neural Network

The construction of our skill distance measure was a three-part process: we began by using a language model to extract a rich semantic representation of the content of job offers, before using a neural network to extract only that content pertaining to skills, allowing us to situate occupations within a high-dimensional space, and finally we calculated the angular distance within this space, which gives us our measure of the skill distance between two occupations.

The principal methodological contribution of this paper lies in our construction of a high-dimensional spatial representation of the skills listed and required by the 532 occupations in the ROME V3 classification, which is pertinent to the quantitative analysis of professional mobility, while also allowing for geometric interpretation. The measure of the skill distance between occupations used in the rest of this article is a direct by-product of this construction. We use FlauBERT (Le *et al.,* 2020), a pre-trained French language model, to analyse the text of job offers.

Our decision to use a neural network to generate our representation is what sets this work apart from the existing literature on the multi-dimensionality of skills. Previous studies have predominantly relied upon methods designed to construct indices or reduce dimensions (principal component analysis, correspondence analysis, etc.), which are mathematically simple to define but yield representations whose pertinence to skills analysis is far from certain – particularly for the purposes of quantitative analysis. We opted for an alternative approach, a supervised method in which all of the pretext tasks used to train the model can be easily explained, are directly connected to the purpose of the representation, and yield a geometrically coherent result. This allows for easier interpretation of our representation, although it does make the construction process more complex.

The fact that our construction of a spatial representation of skills is based on textual data opens the door to certain risks. It is possible that the text of the job offers contains elements unrelated to skills, which are of little use when predicting the associated occupation. The style of writing, the name of the company and the location of the job may all have an influence on our representation. We might also find that, depending on the occupation, job offers tend to demand more skills than the post truly requires, or different skills during different phases of the labour market cycle, as observed by Deming & Kahn (2018). Nevertheless, the size of our corpus of job offers, and the shallow depth of the neural network, allow us to be optimistic that the network will be capable of focusing on those features of the text with the greatest predictive capacity, cutting through the noise to get to the skills descriptions.

1.2.1. Objectives

We thus used a neural network to situate the 532 occupations listed in the ROME classification within a 20-dimensional space, so that the relative positioning of two occupations would provide information as to the likelihood of professional transition from one to the other. The dimensions of this space are not intended to be considered in isolation as measurements of the importance of a specific (or even clearly interpretable) skill for a given occupation. However, the vector formed by these 20 dimensions contains a composite representation of the

4. *To ascertain this value we use the highest level of qualification listed in the Pôle Emploi records for those registered in 2018. This information is derived from the ForCE matching.*
5. *These vertical transitions are constructed using administrative data for the year 2019 (All employees database – job position data -* Base tous salariés, fichier "Postes", *INSEE).*

skills required by a given occupation, and this vector lends itself to interpretation. Once the occupations have been situated in the space, we use the cosine distance between vectors to measure distances between occupations in terms of skills.

Expressed in formal terms, the representation R is the function which associates each occupation in the ROME classification with its corresponding vector in this space.

$$\begin{array}{rccc} R: & \{\text{Codes ROME}\} & \rightarrow & \mathbb{R}^{20} \\ & x & \mapsto & R(x) \end{array}$$

The representation is constructed in such a manner that the geometry of this space can be clearly interpreted as it pertains to skills and professional mobility.

(i) The angle between vectors $R(x)$ and $R(y)$ must reflect the extent to which the skills required by occupations x and y are similar. The cosine of this angle is our measure of skill distance between occupations.

(ii) The norm $\| R(x) \|$ must reflect the expected degree of expertise in the skills required by occupation x.

(iii) The projection of the vector u on the line spanned by $R(x)$ must reflect the degree of expertise in the skills required by occupation x possessed by a person whose skills are represented by u.

(iv) All components of the vector $R(x)$ are positive and represent dimensions of the skills required by occupation x.

The distinction between required skills (direction of the vector) and the degree of expertise in these skills (norm of the vector) is inspired by the questions contained in several of the existing data sets linking skills and occupations, such as the American O*NET classification. In practice, we use the direction of the vectors in the representation of occupations (see below), but their norm also plays a role in the process of constructing that representation.

Figure I illustrates the geometric interpretation of our representation. One consequence of the third property, relating to skills transfer during professional transition, is that the dimension of the space implies that most occupations will be far from each other in terms of the skills they require. In other words, moving from one of these occupations to another would involve starting from zero. We opted to work in 20 dimensions, which means that our representation must cover the skills shared among the 532 occupations in the ROME classification, while also leaving enough flexibility for these skills to be represented in a low-dimensional model.

Our proposed geometric frame has two properties which may appear counter-intuitive. On the one hand, the representation introduces a form of imperfect substitutability between skills. Possessing a very high level of expertise in a given field guarantees a minimum level of ability in every other field which has at least one skill in common. In this respect, the representation does not do a good job of handling situations where transition is totally impossible. Moreover,

Figure I – **Geometric interpretation of the representation *R* of occupations**

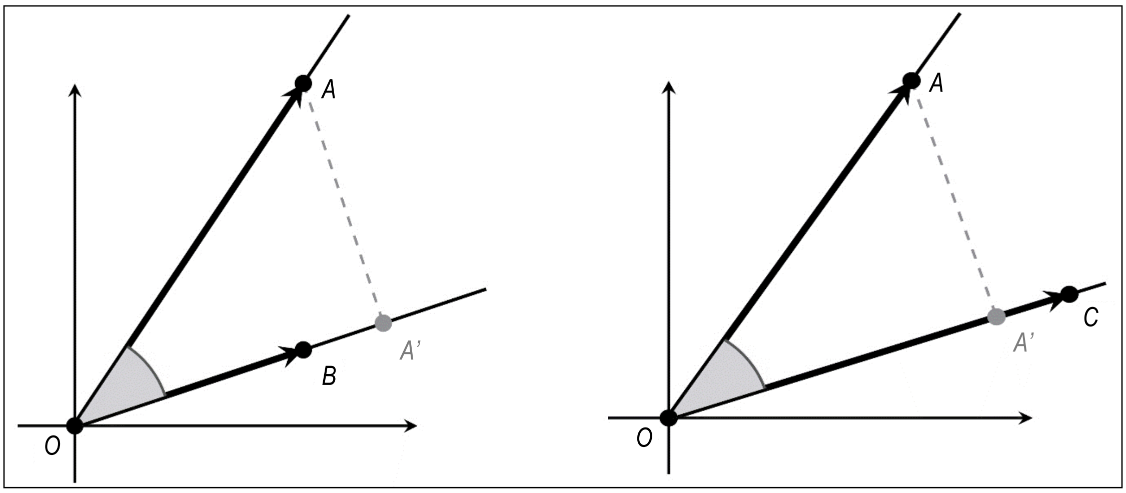

Note: Representation of the transition from an initial occupation A to a new occupation B within the skill space. Our measure of distance between occupations, which reflects the skill gap between occupations A and B, corresponds to the cosine of the angle in grey. Occupation A requires a higher level of skills than Occupation B, hence the distance between the two is greater. If an individual whose skills correspond exactly to those demanded by Occupation A were to transition to Occupation B, his degree of expertise in the skills required by B would be the distance from the origin O to point A. In this case, the degree of expertise is sufficient to make the transition from A to B possible. However, this same individual would not be able to move to Occupation C, which involves the same skills as Occupation B but demands a higher level of expertise.

the binary relation between occupations which indicates that professional transition is possible is not usually transitive; in other words, it may be possible to move from occupation A to occupation B, and from occupation B to occupation C, but never directly from A to C. The underlying idea is that this binary relation reflects the possibility of undertaking professional mobility within a reasonable time frame without receiving specific training. But the sum of two "reasonable" periods of time may no longer be so reasonable.

This geometric representation of the feasibility of professional transitions within the skills space was not used directly in the construction of our measure. Nonetheless, it does illustrate the partial reuse of existing skills when undertaking professional mobility. The further away we move from the initial occupation in terms of the skills required, i.e. the bigger the angle, the smaller the norm of the projection will be, i.e. the less useful existing skills will be. So, for a given level of expertise in the skills required by a new occupation, it is more feasible to make the transition to an occupation with a small angular distance. This highlights the advantage of using angular distance as a measure of the difficulty associated with a professional transition.

In addition to this representation of occupations, our neural network also allows us to assign every job offer to a position within this 20-dimensional space. We use F to represent the corresponding function.

$$\begin{array}{cccc} F: & \{\text{job posting text}\} & \rightarrow & \mathbb{R}^{20} \\ & x & \mapsto & F(x) \end{array}$$

This representation of the texts of our job offer corpus has two advantages:

- It can be used to learn the representation R of occupations. The primary task when training the neural network consisted of comparing the text t of a job offer with an occupation x, by means of their representatives $F(t)$ and $R(x)$, in order to predict the ROME code of the occupation targeted by the job offer.

- Once training was complete, we used the network to conduct a qualitative analysis of the resulting representation. It is also possible to modify the text t of a job offer at the input stage, then observe how the changes affect the representative $F(t)$.

1.2.2. Architecture of the Neural Network

The neural network comprises three blocks, as illustrated in Figure II.

Block 1 is a language model. It takes a textual input and generates a representation of this text in 768 dimensions, encompassing a broad array of semantic problems. We used a pre-trained version of the model, with parameters which remained fixed while training the other Blocks.

Block 2 is a neural network with three layers. It takes as its input the representation of a job offer produced by Block 1, then produces a representation of this offer in 20 dimensions. This step dispenses with the 768 initial dimensions, retaining only the information pertaining to skills and professional mobility. Combined with Block 1, it forms a function F which assigns the text of a job offer to its 20-dimensional representation.

Block 3 is an additional layer which stores the representation R of the 532 occupations in the ROME classification, in 20 dimensions. Comparing the output of Block 2 with the output of Block 3 allows us to assess the performance of the network in its various tasks, and to adjust the training parameters accordingly.

The parameters of Block 1 (dotted background in Figure II) were retrieved from a pre-trained model (FlauBERT. Le *et al.,* 2020), with some slight re-training on a semi-supervised basis, focusing on our corpus of job offers independently of the rest of the model, before freezing these parameters during training of Blocks 2 and 3. The parameters of Blocks 2 and 3 (dark grey layers in Figure II) were random to begin with, then trained on the text of the job offers as explained in the next sub-section.

1.2.3. Training Tasks and Associated Penalties

The representation of occupations R is one of the parameters of the neural network, which evolve as the neural network is trained to perform certain tasks. Choosing appropriate training tasks and associated penalties allows us to impose upon the representation the geometric properties detailed above.

We used the following task-loss pairings to train our neural network.

Predicting the ROME Code Associated With a Job Offer - *WARP Loss*

We decided to adapt the *Weighted Approximate-Rank Pairwise* loss (or WARP) proposed by Weston *et al.* (2011) to the context of our problem. For each text t representing a job offer, we calculated the angle of its representative $F(t)$ with the representative $R(x)$ for each occupation

Figure II – **Diagram of our neural network**

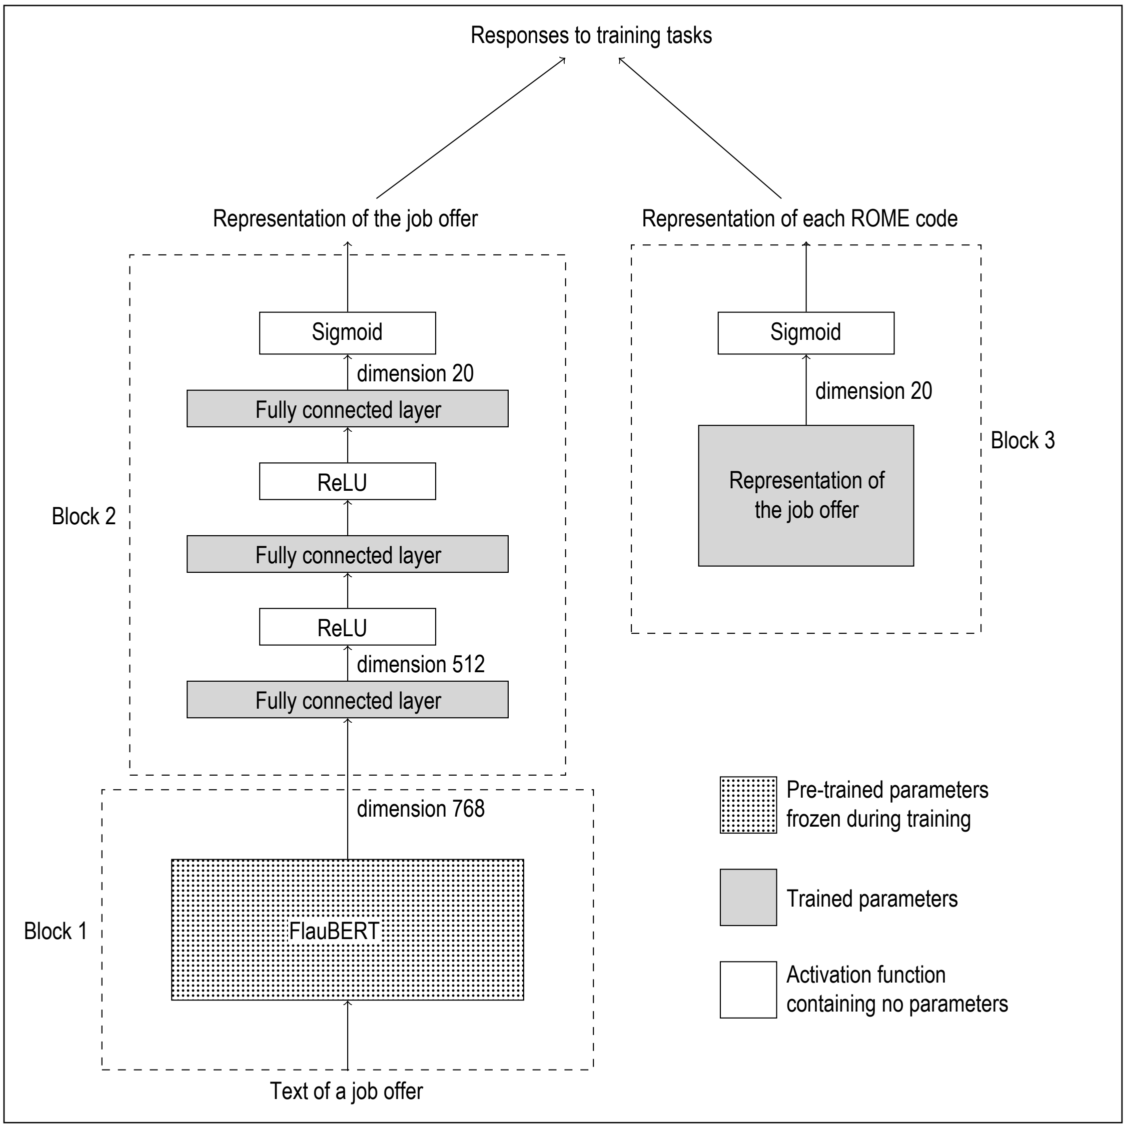

x in the ROME classification. We expected the angle between $F(t)$ and the representative $R(x_0)$ for the occupation corresponding to the job offer to be very small. We thus penalised the network for each occupation x so that the angle between $F(t)$ and $R(x)$ remained smaller than the angle between $F(t)$ and $R(x_0)$.

$$l_{\text{WARP}} = \sum_{\text{ROME codes } x \neq x_0} \left| F(\mathrm{t}) \cdot \left(R(x) - R(x_0) \right) \right|_+$$

where $|\cdot|_+ = max(0,\cdot)$.

Predicting the Mobility Options Suggested in the ROME Classification - Triplet Loss

Triplet loss is a concept whose origins lie in image recognition. If we take three photographs of human faces, with the first two showing the same individual and the third featuring a different person, a facial recognition algorithm should find that the first image bears closer resemblance to the second than to the third.

Similarly, we created triplets of occupations, with an initial occupation x_0, an occupation x which is suggested in the Mobilities section of the ROME classification as a possible career move for individuals in y, and a third occupation which is not suggested as a possible destination for professional transition. We then picked, at random, a job offer corresponding to each of these occupations.

We would expect mobility from x_0 to x to be a better option than moving from x_0 to y. The transition needs to be feasible in terms of the skills required, without too much of a decline in the level of expertise compared with the initial occupation. The following equation can thus be used to determine the extent to which a transition from one occupation to another would be "recommended":

$$d\left(x_0,x\right)=max\left(1-\frac{\|R\left(x\right)\|}{\|R\left(x_0\right)\|},1-\frac{R\left(x_0\right)\cdot R\left(x\right)}{\|R\left(x\right)\|^2}\right)$$

Triplet loss can thus be defined as the difference, with reference to this equation, between the pairs $\left(x_0,x\right)$ and $\left(x_0,y\right)$.

$$l_{Triplet}=|\,d\left(x_0,x\right)-d\left(x_0,y\right)|_+$$

Predicting the Expected Level of Expertise in the Skills Required by an Occupation x - Norm Loss

For each occupation x, we compared the norm of representative $R\left(x\right)$ to a value e_x (normalised between 0 and 1) representing the average qualification level of workers in this occupation.[6]

$$l_{\text{Norm}}=\sum_{\text{ROME codes }x}\left(\|R\left(x\right)\|-e_x\right)^2$$

Predicting Level of Expertise - Vertical Loss

In many sectors, vertical mobility requires candidates to have acquired the skills associated with lower ranks of the same profession. We would thus expect that, for initial-new occupation pairings which are often associated with vertical mobility in the labour market data, the initial occupation would demand fewer skills than the new occupation, in all dimensions encompassed by the notion of skills.

Using $x\prec y$ to denote the set of occupations y for which transitioning from occupation x constitutes a vertical mobility, we define *vertical loss* as

$$l_{\text{Vertical}}=\sum_{x\prec y}\sum_{i=1}^{20}|\,R(x)_i-R(y)_i\,|_+$$

The various penalties associated with these tasks are normalised, passed through the logarithm and added up to calculate a total loss l, which gives us:

$$l=\log\left(1+l_{\text{WARP}}\right)+\log\left(1+l_{\text{Triplet}}\right)+\log\left(1+l_{\text{Norm}}\right)+\log\left(1+l_{\text{Vertical}}\right)$$

1.3. Results

We qualitatively and quantitatively verified the skill distance measures obtained using our spatial representation of occupations. In the appendix to this article, we provide some initial indications of the performance of the algorithm in the pretext tasks, as well as a qualitative validation exercise.

When it came to quantitatively validating our measure, our principal approach consisted of evaluating its capacity to predict professional transitions actually recorded in the labour market.

Our measure is not designed to perfectly predict observed transitions, as career changes may be influenced by any number of factors which have nothing to do with professional skills, such as personal aspirations, gender stereotypes or conditions on the local labour market. Nonetheless, it seems reasonable to assume that there would be more professional transitions between occupations which are relatively close to one another, in terms of the skills involved.

We compared the predictive power of our own measure of distance between occupations with two alternative measures: first, the distance on the graph of the professional mobilities suggested by the ROME classification, and second the cosine distance between the skills vectors associated with initial and new occupations in the ROME classification.

Figure III represents, for each of the three measures of distance between occupations and for each level *k* (between 1 and 10), the number of professional transitions observed from an initial occupation towards the closest *k*-th occupation. This enables us to assess the capacity of each measure to identify the most likely transitions, and to observe the quality of the measures for more distant occupations. We can thus see that our measure of distance (noted Textual distance in the graph) performs better than the distance constructed using the suggestions from the ROME classification (Mobility distance) at 7 of the 10 levels. Furthermore, adding up the transitions explained by the k closest target occupations, our measure is superior for all *k* values from 1 through 10. As for the measure based on the skills defined in the ROME classification V3 (Distance in skills), we can see that it has impressive predictive power for the closest occupations, but is poor at predicting more remote occupations. This is not hard to understand, because the definitions of skills used in the ROME classification V3 are narrow and do not allow for easy generalisation (a shortcoming which has been taken into consideration in the ongoing work to create version 4.0 of the ROME classification). Our measure, however, which was constructed using diverse and unstructured textual data, seems more capable of identifying broad skills categories and making generalisations, as it outperforms its rivals at the highest deciles.

6. *In practice, this level is constructed as the average level of education observed in the occupation, normalised between 0 and 1.*

Figure III – **Number of transitions explained by measures of distance between occupations**

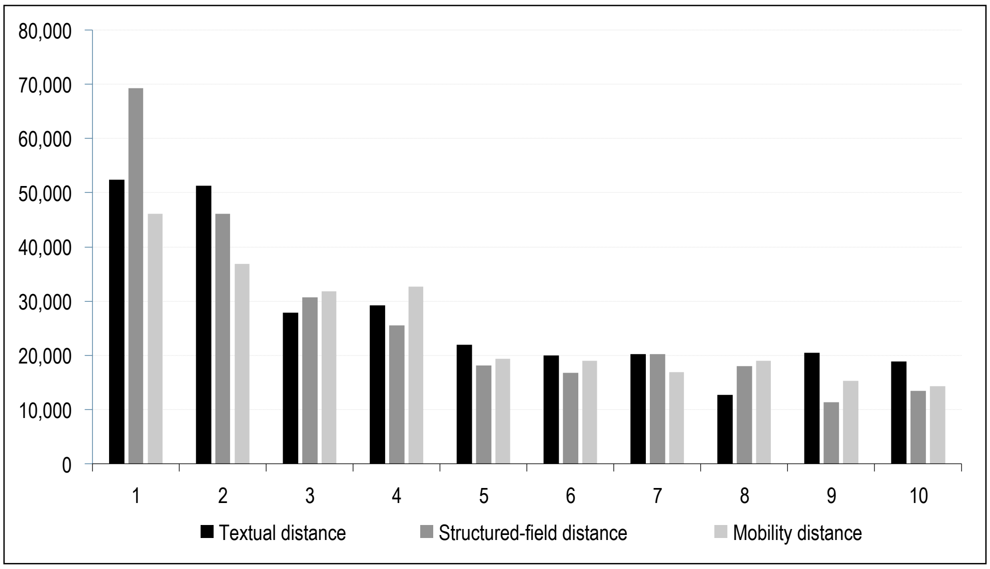

Note: This bar chart represents the number of transitions observed in the labour market in 2019, towards the 10 occupations closest to the initial occupation according to different measures of distance between occupations. We thus compare the textual distance that we constructed to the distances constructed using the Mobilities and Skills structured fields from ROME V3. This diagram makes no distinction between transitions made with or without some form of training.
Source: *Base tous salariés, fichier "Postes"*, 2019, INSEE; authors' calculations.

Table 1 contains the results of our regression analysis, for all pairs of occupations between which at least one transition was observed, of the logarithm for the number of recorded transitions and the different normalised measures of the distance between those occupations. This allows us to compare the capacity of the different measures to predict long distance transitions, i.e. between occupations which are quite distant in terms of the skills they demand, compared with the previous diagram which focused on the most likely transitions. Our measure of distance between occupations does a better job of accounting for variance than the two other measures, and the correlation coefficient is higher. It should come as no surprise that measures of the distance between occupations which focus primarily on skills are only capable of explaining a small proportion of the transitions observed, since they do not take into consideration any of the numerous other factors which may influence professional mobility.

2. Vocational Training, Returning to Work and Career Trajectories

In this section we consider the career progress of jobseekers who undertake training, in terms of their return to employment and their career trajectories in general, comparing their experiences with those of jobseekers who have not undertaken training but are comparable with regard to other observable properties. We begin

Table 1 – **Explanatory power of measures of distance between occupations**

	log(transitions)	log(transitions)	log(transitions)	log(transitions)
Mobility distance	−0.409 (0.0052)			−0.288 (0.0049)
Structured-field distance		−0.4501 (0.0051)		−0.287 (0.0049)
Textual distance			−0.609 (0.0052)	−0.495 (0.0052)
N	149,209	149,209	149,209	149,209
R^2	0.0409	0.0506	0.0935	0.1388

Note: This table contains the results of four regressions for which the dependent variable was the logarithm for the number of transitions between occupations. The first two explanatory variables are distance measures constructed using the Mobilities and Skills structured fields from ROME V3, which we compare with the textual distance constructed for the purpose of this study. Each measure of distance between occupations has been normalised so that the coefficient directly expresses the correlation between the explained and explanatory variables.
Source: *Base tous salariés, fichier "Postes", 2019*, INSEE; authors' calculations.

by detailing the administrative data sources we used and their preliminary processing, then offer a brief explanation of our chosen empirical strategy, before presenting our results.

2.1. Data and Structure of the Sample

We used data derived from the Training, Unemployment and Employment (ForCE)[7] programme, which merges:

- The Historical File on jobseekers (FH, *Fichier historique des demandeurs d'emploi*), containing information on all jobseekers registered with Pôle Emploi in the 10 years preceding the latest edition of ForCE;

- The regional database of participants enrolled in vocational training (Brest, *base régionalisée des stagiaires de la formation professionnelle*), containing details of training courses completed since 2017 by all jobseekers who have at one point undertaken vocational training;

- The manpower movement database (MMO – *Mouvements de main-d'œuvre*, based on data derived from the nominative social declaration, or DSN – *Déclaration sociale nominative*), containing information on employment contracts involving all employees in the private sector since 2017,[8] for all jobseekers present in the HF of the latest edition of ForCE.

We constructed our analytical database using the same broad principles employed by Chabaud *et al.* (2022) in their study of the impact of vocational training on getting jobseekers back into employment, using similar data. We began by creating a database for each month between January 2018 and December 2020, containing all jobseekers registered during that month m (excluding category E / administrative category 5),[9] after first combining any periods of registered unemployment less than 30 days apart. In order to focus our analysis on jobseekers undertaking training for the first time, we used Table P2 in the Historical File to exclude any jobseekers enrolled in training courses in 2017. The sample was matched with the Brest database in order to identify all jobseekers enrolling on a training course for the first time in month m, the group we were interested in for that month – while also excluding the control group of individuals who had received training before month m. We also excluded training programmes regarded as being directly connected to recruitment schemes (POEI, POEC, AFPR), which did not fall within the scope of our study.[10] Finally, matching with the MMO database enabled us to retrieve information regarding the contracts held by jobseekers *(i)* before month m and *(ii)* for each of the 24 months following month m. In so far as this study is devoted to the career trajectories of jobseekers, our principal sample restricts the population of jobseekers to people who have held a stable job at some time in the 12 months preceding month m. We thus identified a reference occupation for each jobseeker which, when linked with the characteristics of the job they found after month m, enabled us to construct the principal variables of interest for this study (skill distance between the two occupations, differential in recruitment difficulties, nature of employment contract, etc.).

We thus retained a large amount of information regarding individuals present in the HF, including their age, gender, marital situation, number of children, level of education and qualification, disability status, residence in urban/rural areas, willingness to undertake geographic mobility, registered unemployed status, reason for registration, nationality, desired occupation, reservation wage, time in unemployment as of month m, total time in unemployment and number of periods of unemployment before this period. These control variables allow us to correct for observable differences between jobseekers who have undergone vocational training and those who have not.

2.2. Empirical Strategy

Let us begin by defining a variable $D_i \in \{0,1\}$ which tells us whether an individual i started training in month m $(D_i = 1)$ or not $(D_i = 0)$. For a variable of interest Y_i (for example, return to employment within 24 months of the month m in which training began), we can define two potential values (using the standard terminology of causal inference) indicating the value of Y_i for an individual i who either did not begin training in the month i (m) or did indeed begin training $(Y_i(0))$. Intuitively, the conditional independence assumption supposes that a treated individual $(D_i = 1)$ would have, in the event that they did not enter training, a similar fate to a control individual i' $(D_{i'} = 0)$ with observable properties X similar to those of i. In formal terms, this can be expressed as follows:

7. For the year 2023T2.

8. Since 2022, information on the current contracts of a large proportion of public sector employees has also been available.

9. Category E / Administrative Category 5 contains all registered jobseekers currently in full time employment, and who are therefore under no obligation to find new employment.

10. The Brest database allows us to distinguish Individual Operational Preparations for Employment (POEI – Préparations Opérationnelles à l'Emploi Individuelles), Collective Operational Preparations for Employment (POEC – Préparations Opérationnelles à l'Emploi Collectives), and Pre-Recruitment Training Actions (AFPR – Actions de Formation Préalables au Recrutement) from the Individual Training Subsidies (AIF – Aides Individuelles de Formation) on which we are focusing.

$$Y_i(0) \perp D_i \mid X.$$

This hypothesis, if verified, enables us to identify the average treatment effect on the treated (ATT), i.e. the average impact of beginning training in month m on the individuals concerned. When interpreting these results, it is important to bear in mind that the treatment we are interested in is the fact of entering training for the first time *during month* m, in comparison with all other scenarios (including no training at all, *or* entering training after month m).

However, it is worth noting that in this case it is probable that those jobseekers choosing to enrol in training are different from untrained individuals in terms of unobservable characteristics (motivation, career ambitions, etc.) likely to be significantly correlated with return to employment and a positive career trajectory in general. This form of bias – which impedes causal interpretation of analyses based on the conditional independence assumption which do not control for differences in observable characteristics – has been extensively discussed in the literature (Lalonde, 1986). It thus seems more prudent to interpret our results as measurements of the existing correlation between vocational training and return to employment, corrected for differences in observable characteristics.

With regard to estimation, there are various ways to control for the differences associated with observable factors X, especially when the dimensionality of X is high. The first is propensity score matching, a method which has been widely used to solve the dimensionality of X problem since it was first proposed by Rosenbaum & Rubin (1983). Our preferred method, known as Double Debiased Machine Learning (DML) (Chernozhukov *et al.*, 2018), relies on the use of nonparametric estimators for the conditional expectation of the variable of interest and the propensity score, which can then be combined to create an estimator capable of withstanding erroneous specification of one of the two terms.[11] Although this solution yields better statistical properties than the classic propensity score matching method, the question of whether or not the DML estimator is biased is just as dependent on the validity of the conditional independence assumption.

2.3. Characteristics of our Sample

In this sub-section we describe the sample of jobseekers used in our analyses. As noted above, the population we are interested in is limited to individuals for whom we can identify a reference occupation and who, as a result, are relatively close to the world of work. As mentioned previously, we also excluded all jobseekers undertaking training directly linked to recruitment programmes (POEI, POEC and AFPR).

Table 2 compares jobseekers known to have held a stable employment contract (permanent post, or fixed-term contract of more than 6 months) within the previous 12 months (the population we set out to study), with those for whom this was not the case. The first sub-population is the only group for which we are able to study the link between training and skill acquisition, as well as the differential in recruitment tension between the initial and subsequent occupations. This sub-population appears to be younger, with a smaller proportion of women and more university graduates; people in this group have been unemployed for less time, and are more likely to undertake training.

Table 3 presents the characteristics of the jobseekers who make up the population of interest to our study, depending on the type of training they have undertaken. We make a distinction between jobseekers not undertaking training and jobseekers taking up training options of 30+ hours, while also considering the type (leading to a diploma, or not) and duration (more or less than 420 hours, equivalent to 3 months of full-time study) of such training programmes.[12]

Compared with those jobseekers who have not enrolled in 30-hour training programmes, the sub-population of jobseekers embarking upon their first 30-hour (or more) training unit are, on average, younger, have been unemployed for less time, and comprise a higher proportion of women and university graduates. These observations are even more salient for those sub-populations of individuals enrolling on courses leading to diplomas, or long training courses involving more than 420 hours of teaching. One more surprising observation is the absence of any notable difference between these populations in terms of career change aspirations. We measured these aspirations by considering the proportion of jobseekers reporting that they were seeking a different occupation from their last stable

---

11. In terms of concrete implementation of this method, packages in Python and R (the language used for this study) are available at: https://docs.doubleml.org/stable/index.html.

12. The number of jobseekers beginning training is subject to major seasonal variation. There is a clear peak in September. We decided to use statistics for individuals enrolling on training programmes between January 2018 and December 2020.

Table 2 – **Descriptive statistics for jobseekers who had (or did not have) stable jobs within the 12 months preceding the month of the study**

	No stable job in the preceding year	With stable job in the preceding year
Female (%)	53.4	48.9
Age (%)		
Under 25	10.2	12.4
Age 25-50	60.1	64.5
Over 50	25.1	18
Level of education (%)		
No high school diploma	50.6	43.2
High school diploma (baccalaureate)	22.1	22.6
Higher than baccalaureate	27.3	34.3
Time in unemployment (months)	31.7	12.6
Training (%)	5.2	8.8
Diploma training (%)	2.1	4.0
Training > 420 hours (%)	2.4	3.7
Observations (thousands)	4,118	908

Source and field: ForCE data, DARES. All jobseekers registered between January 2018 and December 2020, excluding all jobseekers undertaking training directly linked to recruitment programmes (POEI, POEC and AFPR).

Table 3 – **Descriptive statistics by training status for the population of jobseekers who had been in stable employment within the previous year**

	No training	Training	Diploma training	Training > 420 hours
Female (%)	48.9	49.8	51.3	55.7
Age (%)				
Under 25	12.4	12.4	13.3	18.6
Age 25-50	64.4	69.3	71.6	69.3
Over 50	18.1	14.0	10.7	7.7
Level of education (%)				
No high school diploma	43.3	38.5	33.7	32.7
High school diploma (baccalaureate)	22.5	26.5	29.2	31.9
Higher than baccalaureate	34.2	35.0	37.1	35.4
Time in unemployment (months)	12.7	8.3	8.3	8.4
Intended change of occupation (%)	67.5	69.2	69.9	70.8
Distance (when change occurred)	0.51	0.53	0.53	0.54
Observations (thousands)	900	8.0	3.7	3.4

Source and field: ForCE data, DARES. All jobseekers registered between January 2018 and December 2020 after termination of an employment contract (permanent post or fixed-term post of more than 6 months) within the preceding 12 months, excluding all jobseekers undertaking training directly linked to recruitment programmes (POEI, POEC and AFPR).

post.[13] Almost three quarters of the individuals surveyed reported that they were looking for a different occupation from their last post – a proportion which varies little if at all, regardless of treatment status. Moreover, the scale of the skill changes involved – as measured by the skill gap between their previous occupation and the occupation they wish to take up – is comparable for all individuals, whether or not they were beginning training, and irrespective of the nature of that training.

13. For the purposes of this exercise, we used the occupations found in the most detailed version of the FAP classification, containing 225 posts. Using less detailed data, for example FAP 87, does not change the result. Using the FAP classification allowed us to bridge the gap between the ROME classification used in the HF and the PCS system used in the MMOs.

Our identification and estimation strategies were based on the correction of observable differences between control and treated jobseekers, using the DML method (see above). In practice, the control variables used by our algorithm included level of qualification, age, gender, type of contract sought, level of training, experience on the labour market, whether or not jobseekers live in priority urban development zones, marital status, administrative category of Pôle Emploi registration, nationality, the reason for their most recent registration with Pôle Emploi, whether or not they are under obligation to seek employment, the level of their declared reservation wage, their mobility preferences, their desired occupation, the employment zone associated with their place of residence, the number of discrete periods of unemployment and the total time they have been registered with Pôle Emploi over the past ten years.

2.4. Results

All of the results presented in this section were obtained using the *Double Debiased Machine Learning* technique (see above), and are broadly comparable to those obtained using a more familiar propensity score matching method.

Table 4 summarises our results on the correlation between vocational training and return to employment within different time frames (3, 6, 12 and 24 months after beginning training), for all jobseekers enrolling on training courses for the first time between January 2018 and June 2020, and having previously held a stable job in the twelve months before they began training. We also make a distinction between return to employment after training courses leading to diplomas and return to employment in general, before adding a further restriction to include only those jobseekers returning to work in stable jobs (permanent contracts or fixed-term contracts of more than 6 months). As has been extensively documented in the literature (Card *et al.*, 2018), we observed a "lock-in" phenomenon, which is to say a negative correlation between enrolling on a training programme and returning to employment in the short term, due to the time required to complete the training. The correlation between training and return to employment is positive 12 months after beginning training, and remains positive while becoming more pronounced 24 months after beginning training. Our results, which focus on a sub-population relatively close to finding employment, are both qualitatively and quantitatively different from the results reported by Chabaud *et al.* (2022). From a qualitative standpoint, the lock-in effect appears to be more evident and more lasting for our chosen population than it is for the entirety of the population registered with Pôle Emploi. From a quantitative perspective, the corrected differential between trained and untrained jobseekers after 24 months (for our target population) is around 20% less than the result obtained by Chabaud *et al.* (2022) in a study encompassing the entirety of the population registered with Pôle Emploi. These differences may arise from the fact that the individuals in our sample are, by construction, closer to returning to employment than the average jobseeker. It may also be related to the type of training programmes we included in our model. As such, enrolling on a training programme is more likely to significantly reduce return to employment opportunities for this sub-population, for whom employment opportunities are plentiful, even without undertaking further training.

Table 4 shows that while the initial lock-in effect is stronger for training programmes leading to diploma qualifications (which are often long courses of education), the differential between graduates/non-graduates of such programmes after 24 months is greater than the differential for training programmes as a whole. We find something broadly similar when we look at return to stable employment (permanent contracts or fixed-term contracts of more than 6 months). However, the differential between graduates of diploma courses and non-graduates after 24 months is not substantially different from the figure for training courses as a whole. The lock-in effect, however, appears to be more long-lasting: still significant 12 months after training.

Table 5 summarises our principal results on the relationship between training and the career trajectories of jobseekers undertaking various types of training courses (all training programmes and programmes leading to diplomas) 24 months after beginning that training. As explained above, our results concern a sample of jobseekers who had previously been in stable employment within the preceding twelve months. In this table, the dependent variable for return to employment within 24 months is broken down according to the distance between the original and subsequent occupations. We thus distinguish between return to employment in the original occupation ($d = 1$), in an occupation very close to the original occupation ($d \in [2;5]$), in an occupation close to the original occupation ($d \in [6;20]$), and finally in an occupation requiring very different skills from those associated with the original occupation ($d > 20$). We can thus observe that

Table 4 – **Observed differences in return to employment between trained and untrained individuals**

	(1) 3 months	(2) 6 months	(3) 12 months	(4) 24 months
		All types of employment		
All types of training	−0.100 (0.002)	−0.089 (0.002)	0.007 (0.002)	0.069 (0.002)
Diploma training	−0.119 (0.002)	−0.119 (0.003)	0.018 (0.003)	0.086 (0.003)
		Stable employment		
All types of training	−0.063 (0.001)	0.06 (0.002)	−0.011 (0.002)	0.033 (0.002)
Diploma training	−0.078 (0.002)	−0.082 (0.002)	−0.013 (0.003)	0.04 (0.003)

Note: This table contains the results of separate regressions for 4 different dependent variables (in columns, corresponding to return to employment within different time frames) with 2 explanatory variables (the rows, corresponding to the different types of training). The upper section of the table corresponds to return to all types of employment within different time frames, while the lower section repeats these analyses but uses return to stable employment as the only dependent variable (permanent post or fixed-term contract of more than 6 months). Standard errors clustered by occupation × employment zone in parentheses.

vocational training reduces the probability that a jobseeker will return to employment in their original occupation, has a virtually neutral effect on similar occupations, and substantially increases the probability of returning to employment in an occupation requiring different skills from their original occupation. These effects are even more evident when we focus exclusively on training courses leading to diplomas. We can thus say that training in general, and diploma programmes in particular, do appear to coincide with a reallocation of labour supply towards occupations requiring different skills than the posts held before training. This result is striking in so far as it suggests that the entire correlation between return to employment and vocational training involves professional transitions towards very different occupations.

Is the boost in employability which we observed for participants in vocational training backed up by a greater likelihood of moving to a sector of the economy which is looking to recruit? In an attempt to answer that question, Table 6 breaks down the relationship between training and return to employment, looking at whether or not recruitment difficulties are more acute in the market sectors to which trained jobseekers move. The data used to calculate this breakdown are obtained by means of a regression analysis of the return to employment of untrained jobseekers and fixed market effects (employment area × FAP), controlling for the individual characteristics of the jobseekers present in each market. This provides us with an indicator (rate of return to employment for each market) which avoids the familiar measurement problems associated with the non-observable dimensions of the efforts made by jobseekers and businesses.[14] Table 6 shows that, for diploma and non-diploma training programmes alike, the return to employment differential after 24 months does not seem to be driven by more jobseekers switching to

14. DARES use a comparable measurement (labour dispersal by market) to construct a synthetic indicator of tensions on the labour market.

Table 5 – **Differences in return to employment after 24 months for trained and untrained individuals, as a function of the skill distance between their original and new occupations**

	(1) Initial occupation ($d = 1$)	(2) Very similar occupation ($d \in [2;5]$)	(3) Similar occupation ($d \in [6;20]$)	(4) Different occupation ($d > 20$)
All types of training	−0.044 (0.001)	0.005 (0.001)	0.006 (0.001)	0.08 (0.002)
Diploma training	−0.052 (0.002)	0.009 (0.002)	0.005 (0.002)	0.095 (0.003)

Note: This table contains the results of separate regressions for 4 different dependent variables (the columns) with 2 explanatory variables (the rows, corresponding to the different types of training). For example, the coefficients of column (2), Very similar occupation (d∈[2;5]), correspond to the training/no training differential in return to employment (after 24 months) in one of the 4 occupations regarded as being closest to the individual's previous occupation, according to our measure of distance between occupations. Standard errors clustered by occupation × employment zone in parentheses.

Table 6 – **Differences in return to employment after 24 months for trained and untrained individuals, with reference to recruitment conditions in the new occupation**

	(1) Recruitment difficulties less acute than in initial occupation	(2) Recruitment difficulties greater than in initial occupation
All types of training	0.046 (0.002)	0.045 (0.002)
Diploma training	0.050 (0.003)	0.059 (0.003)

Note: This table repeats the analyses from Table 4 but specifies, when constructing the dependent variables, whether individuals returned to unemployment in market sectors where recruitment difficulties were more (or less) acute than in their previous occupations. Standard errors clustered by occupation × employment zone in parentheses.

markets where recruitment shortages are more common than they were in their original occupations. All in all, the impact of training (whether or not it leads to a diploma) on the probability of transitioning towards an occupation where the recruitment demand is stronger appears to be comparable to the probability of transitioning towards an occupation where the recruitment demand is actually weaker. If the aim of the training system is to redress imbalances of supply and demand between different labour markets, this result may appear to be somewhat disappointing. It suggests that refocusing the range of training courses on offer would serve to redirect labour supply towards those sectors experiencing recruitment difficulties; as things currently stand, the impact of training on such transitions appears to be broadly neutral.

* *
*

Does vocational training serve to mitigate structural imbalances in the labour market? This study seeks to provide some form of answer to this question, comparing the career trajectories of jobseekers with and without vocational training. To this end, we constructed an original measure of the skill gap between different occupations, using the text of job offers posted by Pôle Emploi. We used this measure to study the professional transitions undertaken by jobseekers with or without training. For a sample of jobseekers relatively close to finding employment, our results show that, compared with jobseekers without vocational training, trained jobseekers tend to make professional transitions over greater distances, in terms of the skills involved. Our results depend on a conditional independence hypothesis – which is strong in this context – regarding training decisions, and must thus be interpreted with a degree of prudence. In terms of reallocative capacity, the increased likelihood of return to employment after vocational training does not appear to be driven by a surplus of jobseekers switching to occupations where recruitment difficulties are more acute. This result suggests that greater effort to ensure that the range of vocational training options on offer are systematically focused on the skills demanded by occupations in need of manpower would serve to boost the reallocative impact of vocational training. ☐

Link to the Online Appendix:
www.insee.fr/en/statistiques/fichier/8679062/ES547_Frick-et-al_Online-Appendix.pdf

BIBLIOGRAPHY

Baley, I., Figueiredo, A. & Ulbricht, R. (2022). Mismatch Cycles. *Journal of Political Economy*, 130(6). https://doi.org/10.1086/720461

Bana, S. H., Brynjolfsson, E., Rock, D. & Steffen, S. (2019). Job2Vec: Job Title Benchmarking with Collective Multi-View Representation Learning. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19). Association for Computing Machinery, New York, NY, USA, 2763–2771. https://doi.org/10.1145/3357384.3357825

Barnichon, R. & Figura, A. (2015). Labor Market Heterogeneity and the Aggregate Matching Function. *American Economic Journal: Macroeconomics*, 7(4), 222–249. https://doi.org/10.1257/mac.20140116

Basco, S., Liégey, M., Mestieri, M. & Smagghue, G. (2024). The effect of import competition across occupations. *Journal of International Economics*, 153, 104001. https://doi.org/10.1016/j.jinteco.2024.104001

Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis with special reference to education*. New York & London: Columbia University Press for the National Bureau of Economic Research. https://www.nber.org/books-and-chapters/human-capital-theoretical-and-empirical-analysis-special-reference-education-first-edition

Card, D., Kluve, J. & Weber, A. (2018). What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations. *Journal of the European Economic Association*, 16(3), 894–931. https://doi.org/10.1093/jeea/jvx028

Chabaud, M., Bucher, A., Givord, P. & Louvet, A. (2022). Quelles sont les chances de retour à l'emploi après une formation ? DARES, *Document d'études* N° 261. https://dares.travail-emploi.gouv.fr/publication/quelles-sont-les-chances-de-retour-lemploi-apres-une-formation

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68. https://doi.org/10.1111/ectj.12097

Dawson, N., Williams, M. A. & Rizoiu, M. A. (2021). Skill-driven recommendations for job transition pathways. *PLoS ONE*, 16(8), e0254722. https://doi.org/10.1371/journal.pone.0254722

Deming, D. & Kahn, L. B. (2018). Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals. *Journal of Labor Economics*, 36(S1), S337–S369. https://doi.org/10.1086/694106

Eckardt, D. (2022). Are Chemists Good Bankers? Returns to the Match between Training and Occupation. *Working Paper*. https://conference.iza.org/conference_files/TAM_2022/eckardt_d29661.pdf

Fontaine, F. & Rathelot, R. (2022). Le marché du travail français à l'épreuve de la crise sanitaire. *Notes du Conseil d'analyse économique*, 71(2), 1–12. https://www.cairn.info/revue-notes-du-conseil-d-analyse-economique-2022-2-page-1.htm

Gathmann, C. & Schönberg, U. (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics*, 28(1), 1–49. https://doi.org/10.1086/649786

Gentzkow, M., Kelly, B. & Taddy, M. (2019). Text as Data. *Journal of Economic Literature*, 57(3), 535–574. https://doi.org/10.1257/jel.20181020

Guvenen, F., Kuruscu, B., Tanaka, S. & Wiczer, D. (2020). Multidimensional Skill Mismatch. *American Economic Journal: Macroeconomics*, 12(1), 210–244. https://doi.org/10.1257/mac.20160241

Hyman, B. G. (2018). Can Displaced Labor Be Retrained? Evidence from Quasi-Random Assignment to Trade Adjustment Assistance. In: *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, 111, 1–70.

LaLonde, R. J. (1986). Evaluating the Econometric Evaluations of Training Programs with Experimental Data. *The American Economic Review*, 76(4), 604–620. http://www.jstor.org/stable/1806062

Le, H., Vial, L., Frej, J., Segonne, V., Coavoux, M., Lecouteux, B., Allauzen, A., ... & Schwab, D. (2020). FlauBERT: Unsupervised Language Model Pre-training for French. In: *Proceedings of the 12th Language Resources and Evaluation Conference (LREC 2020)*, Marseille. European Language Resources Association, 2479–2485. https://aclanthology.org/2020.lrec-1.302

Lindenlaub, I. & Postel-Vinay, F. (2021). The Worker-Job Surplus. NBER, *Working Paper* 28402. http://www.nber.org/papers/w28402

Marinescu, I. & Rathelot, R. (2018). Mismatch Unemployment and the Geography of Job Search. *American Economic Journal: Macroeconomics*, 10(3), 42–70. https://doi.org/10.1257/mac.20160312

Rosenbaum, P. R. & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55. https://doi.org/10.1093/biomet/70.1.41

Şahin, A., Song, J., Topa, G. & Violante, G. L. (2014). Mismatch Unemployment. *American Economic Review*, 104(11), 3529–3564. https://doi.org/10.1257/aer.104.11.3529

Shaw, K. L. (1984). A formulation of the earnings function using the concept of occupational investment. *Journal of Human Resources*, 19(3), 319–340. https://doi.org/10.2307/145876

Traiberman, S. (2019). Occupations and import competition: Evidence from Denmark. *American Economic Review*, 109(12), 4260–4301. https://doi.org/10.1257/aer.20161925

Weston, J., Bengio, S. & Usunier, N. (2011). WSABIE: Scaling up to large vocabulary image annotation. In: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, 2764–2770. https://dl.acm.org/doi/10.5555/2283696.2283856

APPENDIX

ALTERNATIVE APPROACHES TO EVALUATING THE NEURAL NETWORK

Performance in Pretext Tasks

While training the neural network, we kept an eye on the variation in the general loss function defined above, as well as the capacity of the neural network to correctly predict the ROME code for a specific job offer (i.e. its accuracy).

These two values were calculated throughout the whole training and testing process, with training stopped when we found the parameters which delivered the highest level of accuracy with the training data set – around 80% in our case. We also found that unsupervised retraining of the language model on our corpus of job offers, before training our own neural network, significantly increased (by around 20%) the levels of accuracy we were able to attain. This suggests that the availability of large corpuses of job offers, such as the JOCAS database, can be invaluable when it comes to training models, even if the ROME codes corresponding to these jobs are not known (they can now be imputed by statistical learning techniques).

We observed a general reduction in the level of penalties during the training process. Breaking that down loss by loss, we found that, in spite of the normalisation, the task of predicting a ROME code on the basis of a job offer continued to play the most important role in training. The other losses dropped off fairly quickly, most likely because the constraints involved were easier to satisfy within the geometry of the space.

Qualitative Validation of the Measure of Distance Between Occupations

For the purposes of our qualitative analysis, we focused on the occupation most frequently associated with the jobseekers from the 14 sectors of activity defined in ROME. For each of these 14 reference occupations, we identified the 5 closest occupations according to our criteria, including some which are not listed as suggested mobility options in the ROME classification (see Online Appendices S3 and S4). This second list illustrates the capacity of our neural network to predict plausible professional transitions other than those it was trained on. The results suggest that not only does our measure succeed in identifying the mobility suggestions found in the ROME classification V3, but it also often manages to propose other suggestions which appear to be qualitatively coherent. This corroborates our initial suspicion that the suggestions found in the Mobilities section of the ROME classification are relatively limited in scope, and do not reflect the true range of pertinent professional mobility opportunities. Nevertheless, it is worth noting that our measure of distance between occupations is less consistent for those occupations which are under-represented in our corpus of offers, such as Music and singing (L1202). In the Online Appendix S2, we also introduce a visualization technique to explore the representation that preserves the pairwise cosine distance in the neighborhood of a given occupation while projecting the neighbouring occupations in a 2D space.

Economie et Statistique / Economics and Statistics

Objectifs généraux de la revue

Economie et Statistique / Economics and Statistics publie des articles traitant de tous les phénomènes économiques et sociaux, au niveau micro ou macro, s'appuyant sur les données de la statistique publique ou d'autres sources. Une attention particulière est portée à la qualité de la démarche statistique et à la rigueur des concepts mobilisés dans l'analyse. Pour répondre aux objectifs de la revue, les principaux messages des articles et leurs limites éventuelles doivent être formulés dans des termes accessibles à un public qui n'est pas nécessairement spécialiste du sujet de l'article.

Soumissions

Les manuscrits doivent être adressés au secrétariat de la revue (redaction-ecostat@insee.fr), de préférence en format MS-Word. Il doit s'agir de travaux originaux, qui ne sont pas soumis en parallèle à une autre revue. Les articles peuvent être soumis en français ou en anglais. Le texte d'un article standard fait environ 11 000 mots en français (y compris encadrés, tableaux, figures, annexes et bibliographie, non compris d'éventuelles annexes en ligne). Aucune proposition initiale de plus de 12 500 mots (11 500 mots pour les soumissions en anglais) ne sera examinée.

La soumission doit comporter deux fichiers distincts :

- Un fichier d'une page indiquant : le titre de l'article ; les prénom, nom, affiliations (maximum deux) et adresses e-mail et postale de chaque auteur ; un résumé de 160 mots maximum (140 mots pour les soumissions en anglais) qui doit présenter très brièvement la problématique, indiquer la source et donner les principaux axes et conclusions de la recherche ; les codes JEL et quelques mots-clés ; d'éventuels remerciements.
- Un fichier anonymisé du manuscrit complet (texte, illustrations, bibliographie, éventuelles annexes) indiquant en première page uniquement le titre, le résumé, les codes JEL et les mots-clés.

Les propositions retenues sont évaluées par deux à trois rapporteurs (procédure en « double-aveugle »). Les articles acceptés pour publication devront être mis en forme suivant les consignes aux auteurs (accessibles sur https://www.insee.fr/fr/information/2410168). Ils pourront faire l'objet d'un travail éditorial visant à améliorer leur lisibilité et leur présentation formelle.

Publication

Les articles sont publiés en français dans l'édition papier et simultanément en français et en anglais dans l'édition électronique. Celle-ci est disponible, en accès libre, sur le site de l'Insee, le jour même de la publication ; cette mise en ligne immédiate et gratuite donne aux articles une grande visibilité. La revue est par ailleurs accessible sur le portail francophone Persée, et référencée sur le site international Repec et dans la base EconLit.

Main objectives of the journal

Economie et Statistique / Economics and Statistics publishes articles covering any micro- or macro- economic or sociological topic, either using data from public statistics or other sources. Particular attention is paid to rigor in the statistical approach and clarity in the concepts and analyses. In order to meet the journal aims, the main conclusions of the articles, as well as possible limitations, should be written to be accessible to an audience not necessarily specialist of the topic.

Submissions

Manuscripts should be sent to the editorial team (redaction-ecostat@insee.fr), preferably in MS-Word format. The manuscript must be original work and not submitted at the same time to any other journal. It can be submitted either in French or in English. The standard length of an article is of about 10,000 words (including boxes if any, tables, figures, appendices, bibliography, but not counting online appendices if any). Manuscripts of more than 11,500 words will not be considered. Submissions must include two separate files:

- A one-page file providing: the title of the article; the first name, name, affiliation-s (at most two), e-mail et postal addresses of each author; an abstract of maximum 140 words, briefly presenting the research question, data and methodology, and the main conclusions; JEL codes and a few keywords; acknowledgements if any.
- An anonymised manuscript (including the main text, illustrations, bibliography and appendices if any), mentioning only the title, abstract, JEL codes and keywords on the front page.

Proposals that meet the journal objectives are reviewed by two to three referees ("double-blind" review). The articles accepted for publication will have to be presented according to the guidelines for authors (available at https://www.insee.fr/en/information/2591257). They may be subject to editorial work aimed at improving their readability and formal presentation.

Publication

The articles are published in French in the printed edition, and simultaneously in French and in English in the online edition. The online issue is available, in open access, on the Insee website the day of its publication; this immediate and free online availability gives the articles a high visibility. The journal is also available online on the French portal Persée, and indexed in Repec and EconLit.

N° 547 - 2025

Economie et Statistique

Economics and Statistics

ISBN 978-2-11-162505-1 - ISSN 0336-1454 - ECO 547
Parution décembre 2025 - PRIX : 22 €