WORKING PAPER | ISSUE 03/2023 | 13 MARCH 2023

AI ADOPTION IN THE PUBLIC SECTOR: A CASE STUDY

LAURA NURSKI

This case study illustrates the drivers of and barriers to artificial intelligence adoption by organisations, and acceptance of AI by workers in the public sector. Several factors were crucial in the successful adoption of a human-centred approach to AI, including a fast discovery phase that involved workers (or end users) in the development early on, and aligning human resources, information technology and business processes. Subsidy support mechanisms were also specifically targeted and acquired to support the adoption.

However, making AI support available to workers proved insufficient to ensure its widespread usage throughout the organisation. The slow adaptation of existing work processes and legacy IT systems was a barrier to the optimal usage of the technology. Moreover, the usefulness of the technology depended on both the task routineness and worker experience, thereby necessitating a rethinking of the work division between technology and workers, and between junior and senior workers.
Successful human-centred roll-out of AI in Europe will therefore depend on the availability of, or investments in, complementary intangible organisational capital. Very little is currently known about these investments.

The author is grateful to Tom Schraepen (Bruegel) for research assistance, to Mia Hoffmann (Georgetown's Center for Security and Emerging Technology) for comments on earlier versions, and to the contacts at the case organisations, who provided their cooperation and input to the study.

Laura Nurski (laura.nurski@bruegel.org) is a Research Fellow at Bruegel

bruegel

Recommended citation:
Nurski, L. (2023) 'AI adoption in the public sector: a case study', *Working Paper* 03/2023, Bruegel

Table of contents

1 Introduction

1.1 Productivity and technology acceptance

Artificial intelligence (AI) is a new general purpose technology (GPT) expected to bring productivity gains, but also to impact the nature and quality of work. The introduction of previous GPTs, including electricity and computers, has shown a long lag between the adoption of new technologies in the production of goods and services, and widescale observable increases in total factor productivity. This puzzling observation is often dubbed the *"productivity paradox"* (Landauer, 1995), ie the phenomenon that large investments in new IT technologies by firms have not been accompanied by subsequent increases in national productivity statistics. Two crucial factors in explaining this paradox are the slow adaptation of business processes and the underutilisation of technology by workers (Devaraj and Kohli, 2003).

When organisational processes and incentives are not aligned with technology use, workers will not fully adopt new technologies in their work (Atkin *et al*, 2017). AI applications equally suffer from such underutilisation, as studies of the banking (Xu and Zhu, 2021), retail (Kawaguchi, 2020) and healthcare (Jauk *et al*, 2021) sectors demonstrate. Furthermore, it takes time, money and willingness for business models and organisational models to adapt to the use of new technology. Such adaptation requires investment in people practices or human resources (HR) practices, including training, performance evaluations and hiring (Bloom *et al*, 2012). Investment is also needed in organisational practices, such as business process reengineering, decentralisation or organisation redesign (Bresnahan *et al*, 2002).

To analyse the drivers of and barriers to AI adoption by organisations, and acceptance of AI by workers, we investigate a specific case in this study. The study does not serve as a 'star case' (which demonstrates reproducible best practices) or as a 'research case' (which aims to improve economic theories of causality) (Baker and Gil, 2013). Instead, this is a 'teaching case' that illustrates scientific theories in a practical example and bridges several disciplines (including IT, management, organisational behaviour, psychology and economics) by linking theories from respective domains in one case. The goal is to identify pitfalls in the process of technology adoption and to provide some lessons for both policy and business. This case study is part of the Future of Work & Inclusive Growth project[1] at Bruegel, which aims to identify the impact of technology on the nature, quantity, and quality of work.

1.2 The organisations in this case study

We analyse AI adoption by Flanders Investment and Trade, a public organisation, which was assisted by Radix, a private firm[2]. Note that throughout the paper we refer to a list of case-study materials through numerals shown in square brackets. The annex lists the case-study materials.

Flanders Investment and Trade (FIT) is the trade promotion organisation (TPO) of Flanders, a region of Belgium. TPOs are facilitative agencies that promote and stimulate trade by providing information, linkages, technical advice, marketing and policy advocacy (Giovannucci, 2004). Their activities can be grouped into four broad categories: product and market identification and development; trade information services; specialised support services; and promotional activities abroad (Jaramillo, 1992). FIT's mission is to internationalise the economy of Flanders by assisting Flanders-based companies in their export effort ('trade') and by attracting foreign companies and investment to the region ('invest'). Alongside delivering trade and investment services, FIT engages in promotional and development activities including the hosting of events and publication of market insights. FIT has six regional offices in Flanders and Brussels (employing about 150 people) and 100 local offices abroad (employing about 180 people).

Radix is a Belgian AI solution provider, founded in 2018. It has a team of 40 engineers and solution leads across two offices in Flanders and Brussels. Radix provides a portfolio of AI solutions to improve operations in a range of industries, including manufacturing, transportation, financial services and the public sector.

1.3 Selection of the case

The case was found through the website of the AI developer (Radix), which showcases client stories. Several Radix client stories were relevant for the Future of Work and were therefore considered. Among them were two clients in the human resources and public employment sectors: an AI-supported orientation test developed for the Flemish public employment agency, and an AI-powered job-matching algorithm developed for a private HR services company. AI will likely play a major role in matching job seekers to job vacancies in the labour markets of the future. Both the opportunities and

[1] See https://www.bruegel.org/future-work/future-work-and-inclusive-growth-europe.
[2] See https://welcome.flandersinvestmentandtrade.com/ and https://radix.ai/.

potential dangers of this application are currently being studied and debated widely, with specific focus on the risk of increasing discrimination in the labour market. However, in this particular case study, the goal is to study AI not in the job-matching process, but in the production process itself.

FIT was highlighted as a Radix client that adopted AI in one of their core business activities: answering trade-related questions from Flemish companies looking to do trade abroad. Other client cases of this AI developer with applications in the production process included: a production planning algorithm that improves on-time-delivery of production orders, taking less time than a human planner; an algorithm that improves vaccine development by counting and reporting colony forming units; and an algorithm that automatically tags new articles of a news supplier with topical hashtags. We selected the question-answering algorithm for FIT over these other examples because it fitted the current narrative of AI replacing routine cognitive tasks of knowledge workers. Another reason was that the developer noted in their FIT client profile both productivity increases (27 percent time savings, 36 percent more questions answered) and job satisfaction improvements (focus on more complex cases and other parts of their jobs) [8 – see the annex], which fitted our goal of studying both productivity and job-quality effects.

1.4 Methodology

The case was studied through the collection and analysis of several data sources. First, desk research was performed on the existing scientific theories and evidence on AI adoption and acceptance. This desk research resulted in the publication of several blog posts and papers on these topics (see for example Hoffman and Nurski, 2021a, 2021b). Second, desk research was carried out on publicly available information on the cases, most notably the respective websites of FIT and Radix. In a third step, interview guides were developed on the topics on technology adoption and acceptance for several interviewee targets. Interviews were conducted with FIT's AI lead and HR lead (see [4], [10], [13]) and with four 'end users' of one specific AI application at FIT, also known as 'case handlers' (see [12]). The four end users (two men and two women) were stationed in four different offices: France, Germany, Italy and the USA. Depending on the internal organisation of the office, some of the interviewees specialised in certain regions of their country, while others specialised in certain industries in that country. A final data source consisted of collected documents, including slide decks, screenshots and training materials. The full list of case study materials can be found in the annex.

2 AI adoption by the organisation

2.1 Adoption process

2.1.1 Timeline

As part of its digital innovation strategy (see section 2.2.1), FIT is adopting AI across a range of activities in its primary services, namely the trade and invest services. Over four years (2017 to 2021), FIT went through three AI project cycles to: (1) experiment with proof-of-concepts (POCs), (2) build an AI strategy, and (3) set-up the necessary data infrastructure.

Table 1: Summary of phases in the AI adoption process

Year	Phase	Goal
2017-2019	AI proof-of-concepts	Quick POCs to experiment, learn and discover opportunities
2020	AI strategy	Assessing current as-is AI maturity and developing a roadmap towards the desired to-be state of AI adoption
2020-2021	Data infrastructure	Install required infrastructure for centralising and processing all internal and external data sources.

Source: Bruegel based on [4].

2.1.2 Phase 1: Developing AI proof-of-concepts (2017-2019)

In the first phase, FIT familiarised itself with AI technology to discover opportunities and investigate whether it would be useful to explore further. An external AI agency (Radix) set up a 'fast discovery workshop' for FIT's AI lead to screen FIT's business processes for potential AI opportunities [7]. This workshop consisted of a series of brainstorming exercises between the AI developer and the organisation looking to adopt AI. First a longlist of ideas was assembled by gathering ideas from different stakeholders; next the ideas were analysed and prioritised in light of their technical feasibility and business value; finally, effort and value estimations were made for the selected opportunities [14].

This process generated five proof-of-concepts (POCs) for using AI to support FIT's core business services, namely the trade and invest services. They ranged from information gathering on foreign companies through web scraping, lead detection of potential clients through social listening, and predictive modelling for marketing based on likelihoods to invest and trade [4]. This list of opportunities was prioritised according to their business value and technical feasibility (effort and complexity of implementation) (see Figure 1). The POC that came out as a 'quick win' (high value, low complexity) was a question-answering algorithm for FIT's 'trade cases'[3], aimed at partly automating the process of answering trade questions from Flemish companies about foreign markets. Using natural language processing, trained on a large dataset of past trade questions and answers, the algorithm was designed to retrieve past answers to frequently asked routine questions. The 'high value' was estimated because of the large share this task takes up in the workload of case handlers (namely, 60 percent to 70 percent of their workload). The 'low complexity' was estimated due to the availability of high quality 'off the shelf' natural language processing (NLP) models that could be trained on FIT's large history of five years of previously answered questions (about 10,000 per year). Finally, an algorithm was designed to retrieve past answers to routine questions, so that FIT advisors could spend more time on the complex questions. The application acts as an AI-powered search engine, not just comparing individual words, but interpreting the entire body of the question and finding the most relevant past answer.

Figure 1: Value-complexity matrix for prioritising AI opportunities

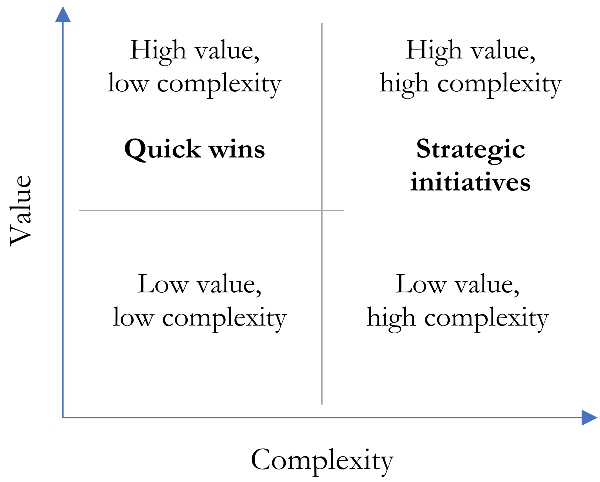

Source: [7].

The 'trade cases' question-answering POC was further developed into a complete AI product by integrating the algorithms' recommendations into FIT's existing Customer Relationship Management software (CRM), Microsoft Dynamics. To evaluate and improve the quality of this first minimum viable product (MVP), the developer conducted 10 interviews across several of FIT's international offices and assessed the results for 175 new trade questions that were handed to the AI. In each case, the algorithm suggested five previous answers, meaning about 875 AI-suggested answers were evaluated.

[3] A 'trade case' is a question from a Flemish company about a foreign market, that concerns services of FIT, for example inquiries about the size or customs of a local market, potential foreign business partners, trade regulations or barriers, subsidies, or market opportunities. See section 3.1 below for more detail on the business process and AI support.

The developer used staff members' personal memories of past cases by asking them if a better answer from the past existed, and then analysed why the algorithm did not retrieve the most relevant answer. Just as workers learn how to improve their answers over time, the algorithm was retrained based on the corrections of FIT staff. Reasons for missing better answers from the past included: unrecognised synonyms (same topic but different words), wrong language (same topic but different language, eg English, Dutch or other language), unclear link (same topic but not explicitly mentioned), wrong focus (AI didn't focus on right words), and out-of-vocabulary (AI didn't know certain words). By taking into account staff member feedback, the hit rate (cases in which the AI found a relevant answer to a question) increased from 51 percent to 62 percent [7]. Involving users in the design of the algorithm thus improved its quality (and therefore useability, see 3.3.2) substantially.

2.1.3 Phase 2: Building an AI strategy (2020)

The first phase showed that it was possible and opportune to expand the adoption of AI in a wider range of FIT's processes. In the second phase, they took a step back from the original five POCs and took a more structural approach to AI by building an AI vision and strategy (or AI 'blueprint') for the organisation. With the help of three external experts, an AI maturity assessment was done, followed by the design of a future vision and a roadmap to move from the as-is situation to the desired to-be state [4].

The methodology for building the AI strategy consisted of three building blocks. First, an enterprise architecture was drawn up, mapping the current business processes on applications, data layers and technical systems. Second, an AI maturity assessment was conducted to assess the 'as-is' state of AI maturity and to develop an AI roadmap of potential 'to-be' states of AI adoption. The third part of the AI strategy related to training and human resources. It included setting up an AI unit responsible for AI impact and dissemination at FIT, training everyone at FIT on basic AI literacy, and specific training for the digital marketing team on data-driven marketing strategies and tools.

The external experts classified the as-is state of FIT's AI maturity at 'AI ready', which is the second level of maturity in their assessment framework:

- AI Novice: AI novices have not taken proactive steps on the AI journey and, at best, are in assessment mode.
- AI Ready: Sufficiently prepared to implement AI in terms of strategy, organisational set-up and data availability.

- AI Proficient: A reasonable degree of practical experience and understanding of how to move forward with AI. There are still gaps and limitations.
- AI Advanced: A good level of AI expertise and experience, with a proven track record across a range of use cases. Good operational procedures in place.

The AI roadmap towards the to-be state was drawn up to move through three states. In a first stage, FIT would use self-service business analytics[4] and dashboarding apps (such as Power BI and Azure data services) and ready-made AI supported insights (for example Office 365 workplace analytics) to build a data foundation and support a data-driven decision-making culture. In a second stage, FIT could use solution-specific AI services and AI-based content understanding (for example chatbots and Application Programming Interfaces (APIs) to Natural Language Processing (NLP) models) to build an FIT conversational knowledge platform. Finally, in the third stage, FIT could adopt advanced cloud infrastructures and open machine-learning frameworks, as well as develop their own custom data science and deep AI capabilities to support the digital marketing pipeline (for example on targeted ads, leads and direct marketing).

2.1.4 Phase 3: Setting up the necessary data infrastructure (2020-2021)

From the assessment in phase 2, it became clear that FIT lacked the required infrastructure for large-scale AI projects that, for example, require the processing of unstructured data in real time. The first step in the roadmap therefore consisted of building a data hub (or data vault) for absorbing data from different internal data sources [4]. These internal sources included FIT's accounting system, Enterprise Resource Planning (ERP) system, website, CRM system and two old legacy systems that still fed into the CRM. The data hub would also centralise and ingest all purchases of external data, like company databases. On top of the physical infrastructure for storing data, an operational database layer would be built around customers, products, accounts and transactions. This data layer would feed into an API access layer that would grant different business applications access to and monitor their use of the data. This set-up would serve as the basis for all future data consumption (both structured and unstructured), data sharing and exchange, data monitoring and access management. By supporting near real-time data processing and reporting, it would serve as the foundation for all future AI development.

---

[4] Self-service analytics is a form of business intelligence (BI) in which line-of-business professionals are enabled and encouraged to perform queries and generate reports on their own, with nominal IT support. (https://www.gartner.com/en/information-technology/glossary/self-service-analytics).

2.2 Drivers and barriers to adoption

An organisation's decision to adopt a new technology is influenced by the technological, organisational and environmental context (Baker, 2012; Hoffmann and Nurski, 2021). According to a Europe-wide company survey (European Commission, 2020), the main reasons for firms to not adopt AI are a lack of financial means, human capital and data availability, both within the firm and from the external environment (Hoffman and Nurski 2021). Table 2 lists drivers and barriers that were identified in this case study in each of the three contexts, while the following paragraphs dive deeper into each of the factors.

Table 2: Identified drivers and barriers to AI adoption at FIT in the technological, organisational and environmental context

	Identified drivers & facilitators	Identified (overcome) barriers
Technological context	Expected productivity gains Data availability High trialability	Lack of compatible IT infrastructure
Organisational context	Leadership and management support	
Environmental context	Competitive pressures	Access to skilled labour and external funding

Source: Bruegel based on Baker (2012), interviews, documents and websites (see the annex).

2.2.1 Main driver of adoption: competitive environment

As a small, open economy, international business is a key factor in the economic development of Flanders. In 2021, Flanders imported €378.8 billion worth of goods and services and exported €380.5 billion, putting the Flanders region in the top 20 of global exporter countries (WTO Stats dashboard). Top exported products include pharmaceutical, chemical, and mineral products, and machinery, electronic and transport equipment. The main trading partners are neighbouring countries Germany, France and the Netherlands, and intra-EU trade represents two-thirds of total exports from Flanders [2]. While separate numbers are unavailable for Flanders, export from Belgium as a whole supports 843,900 jobs in Belgium out of five million total employment (Rueda-Cantuche *et al*, 2021).

TPOs around the world compete for local investments by multinational companies and need sophisticated approaches to attract, and keep foreign investors (Zanatta *et al*, 2006). FIT considers digitalisation a key factor in its strategy to stay competitive in this international landscape [3]. FIT therefore aims to be an 'early adopter' (Rogers, 1983) in digitalisation. The achievement of this goal is recognised by its environment, as FIT is considered one of the best practices for digitalisation and AI adoption by the European Commission [4 and 9].

The digitisation of FIT reflects the wider digital transformation of the Flemish government and the Flemish Digital Strategy, building on the Flemish Data Strategy that was approved on 18 March 2022 [5]. While the digital strategy is still being built, the Flemish government aims to reach a top-five spot in the European ranking of digital public services, as measured by the Digital Economy and Society Index (DESI) [6].

2.2.2 Overcoming financial barriers: external financing

For each of the three phases, external project subsidies were acquired for the specific goal of digitalisation and AI adoption, either directly or indirectly financed by public funds. The first stage (AI POCs) and third stage (data infrastructure) took place within the framework of Flanders Accelerates, which is FIT's internationalisation strategy for the Flemish economy. The execution of this strategy is supported by a combination of European and regional (Flemish) funds. For the period 2017-2022, FIT received €1.8 million from the European Regional Development Fund (ERDF) and €1.6 million from the Fund for Accompanying Economic and Innovation Policy (Hermes Fund), managed by the Flemish Innovation and Entrepreneurship agency (VLAIO). Both funds were awarded specifically for FIT's digitalisation strategy.

The second phase (building the AI strategy) was specifically and directly supported by the Structural Reform Support Programme (SRSP), managed by the European Commission's Directorate-General for Structural Reform Support (DG Reform), the EU body that helps countries design and implement reforms as part of their efforts to support job creation and sustainable growth. The Commission provided support over a 12-month period in the form of technical advisory services by entities with substantial experience in the development of blueprints for AI for public administrations [9]. The advisory services supported the three elements of the AI strategy discussed above, namely: (1) developing an AI maturity assessment; (2) recommending a future architecture and roadmap for AI deployment; (3) proposing curricula for AI-related training of FIT staff. DG reform features the project on its website as inspiration for other EU countries [9].

2.2.3 Overcoming human and organisational barriers: hiring and training

Following FIT's digitalisation and innovation strategy (see 2.2.1) the management team decided that *"FIT wanted to join the AI train"* [10]. A business and information systems engineer with seven years' experience in IT was then hired as a project manager data architecture & artificial intelligence in 2017 – referred to in this case study as the 'AI lead'. Management support was also made public when both FIT's CEO and the Head of IT, personnel and finance endorsed the AI adoption, digitalisation and data-driven decision-making of FIT on the occasion of DG Reform's spotlight on the project (15 December 2020 [11]).

The experts' AI strategy and roadmap (see 2.1.3) recommended setting up an AI-specific unit responsible for AI impact and dissemination. This unit would be in charge of developing a general AI terminology to be used in FIT and further developing the AI strategy. It would also build in-house knowledge of workflows for machine learning (ML), data science (DS) and AI projects. To achieve this, it would include employees with mathematical and statistical backgrounds or experience.

Besides setting up the dedicated AI unit, the roadmap also included training curricula for specific groups within FIT's organisation. All employees at FIT would receive training in order to be ready to welcome and use AI. This training includes understanding the value of data, understanding the impact of AI on business operations, and mastering a general AI vocabulary. The digital marketing team would receive a targeted training on the usage of data to boost the business. This training included, beyond the AI basics, also working with data driven marketing strategies and tools like Google Analytics.

2.2.4 Overcoming technical barriers: data availability and IT compatibility

In the discovery phase (phase 1, see 2.1.2) the AI lead and the AI agency (Radix) scouted opportunities for using AI in the FIT organisation. The goal was to build a list of potential business cases in order to develop proofs-of-concept. This fast-discovery phase facilitated trialability, ie the ability to experiment with an innovation before commitment, which reduces uncertainty and facilitates adoption (Lundblad, 2003). The AI lead based his search on two criteria: (1) potential productivity gains and (2) data availability. To assess potential productivity gains, he asked FIT employees which tasks they currently spend a lot of time on, hindering them in their work, or what tasks they could be supported with. Among one group of employees (the case handlers, see 3.1 for more details), answering repetitive trade questions from Flemish companies looking to trade or invest abroad featured consistently among the top answers. This process – internally known as 'trade cases' – was also the business process that had the most historical data, which was ultimately, the main determining factor in choosing the POCs of intelligent decision-making systems. The current CRM system (Microsoft Dynamics) kept track of all past trade questions with their respective answers. Even the cases that were originally stored in the legacy system (Lotus Notes) were imported into the new CRM [10], meaning that a very large history of answered questions were available on which to train the AI model[5]. Research shows that firms are indeed more likely to build AI on top of existing data-driven applications than to invest in completely new applications (Hoffmann and Nurski, 2021).

From the experts' assessment in phase 2, FIT learned that it lacked the required infrastructure for large-scale AI projects that, for example, require processing unstructured data in real time. Indeed, technological readiness and existing digitalisation is especially important for AI adoption, since digital technologies are hierarchical, meaning the use of AI systems requires other 'lower' technologies such as data storage and computing power (Zolas *et al*, 2020). Without a way to collect, store, move and transform data, companies cannot begin to learn from their data or use it to support intelligent decision making (Figure 2; Hoffmann and Nurski, 2021). Phase 3 therefore consisted of building the recommended IT architecture that could support the AI roadmap, including a physical data-storage infrastructure, an operational data layer and an API access layer.

Figure 2: The hierarchical nature of digital technologies

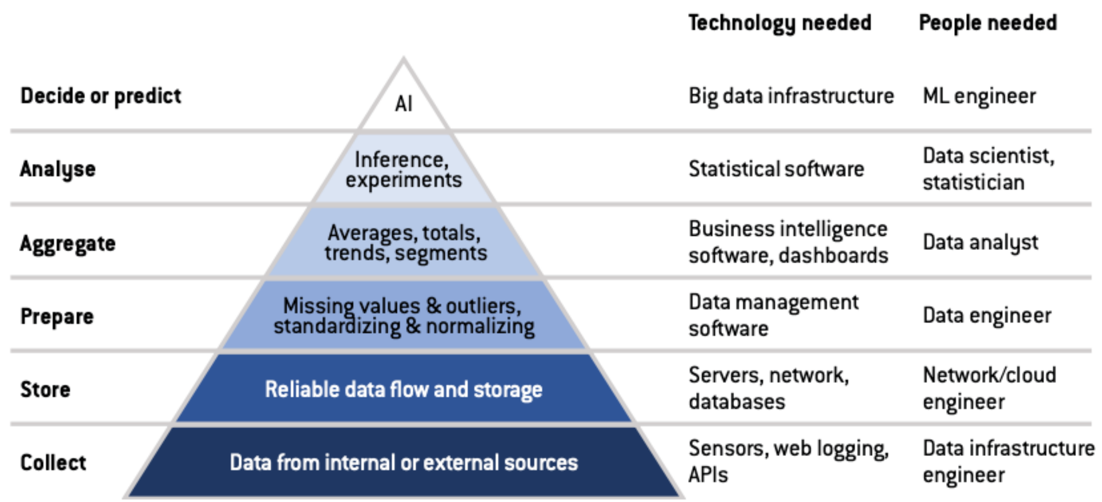

Source: Bruegel based on Monica Rogati, 'The AI hierarchy of needs', *Hackernoon*, 12 June 2017, available at https://hackernoon.com/the-ai-hierarchy-of-needs-18f111fcc007.

[5] Ultimately, only the most recent years of training data were used to ensure that the provided answers were not outdated.

3 AI acceptance by staff members

To assess if, how and to which extent staff members accepted and used AI algorithms in their daily activities, section 3.1 zooms into one of the original POCs described in section 2.1.2, namely the AI assisted question answering of 'trade cases'. Sections 3.2 and 3.3 are based on interviews with four end users of this application, also known as case handlers. The four end users (two men and two women) were stationed in four different offices: France, Germany, Italy and the USA. Depending on the internal organisation of the office, some of the interviewees specialised in certain regions of their country, while others specialised in certain industries of that country (see 1.4 and [12]).

3.1 Studied algorithm: AI-assisted question answering

Answering 'trade cases' is one of the core business activities of FIT. A 'trade case' is a question from a Flemish company about a foreign market, that concerns the services of FIT, for example inquiries about the size or customs of a local market, potential foreign business partners, trade regulations or barriers, subsidies or market opportunities. In 2021, FIT trade officers made 11,152 such tailor-made export recommendations to Flemish companies [1]. A typical FIT office abroad handles about 200 cases per year, meaning about four cases every week. These questions range from very routine information requests (such as providing a list of 10 accountants in Paris) to very non-routine recommendations (such as helping to choose the next export market for an expanding Flemish company).

The existing digital tools (before the AI adoption) for answering trade cases consisted of Microsoft Dynamics as a CRM software, used for receiving, assigning and answering incoming questions, and Microsoft SharePoint as a collaborative document management and storage system, used for storing relevant information on local markets, partners and regulations, and for storing documents drawn up for previous cases. Therefore, trade cases are received and answered through the CRM system (either directly as e-mails or manually inputted from phone calls), but the information and knowledge needed and used to answer them is stored in SharePoint.

The studied algorithm is aimed at partly automating the process of answering trade cases. Using natural language processing, trained on a large dataset of past trade questions and answers, the algorithm was designed to retrieve past answers to frequently-asked routine questions. It acts as an AI-powered search engine, not just comparing individual words, but interpreting the entire body of the question to find the most relevant answer (see also section 2.1.2 for more information on the development of the AI). The algorithm thus retrieves previous answers that might be relevant to the current incoming question and shows the text of the answer and other relevant information from the previous case, such as the name, sector and country of the firm. The AI suggestions were integrated into the CRM system as an extra tab on the CRM screen displaying the information on the incoming question.

3.2 Framework for user acceptance and actual use

To analyse user acceptance of the AI at FIT, we used the Unified Theory of Acceptance and Use of Technology model (see Figure 3 and Venkatesh *et al*, 2003) from the information systems literature[6]. In this model, actual uptake of new technologies is driven by a user's intention to use, and facilitating conditions for use. Facilitating conditions include adequate support infrastructure and assistance. The behavioural intention to use technology is called the user 'acceptance'. This acceptance is shaped by how the user perceives the technology's usefulness and its ease of use and any influence of the social environment of the user. Perceived usefulness is defined as *"the degree to which a person believes that using a particular system would enhance his or her job performance"* and perceived ease-of-use is defined as *"the degree to which a person believes that using a particular system would be free from effort"* (Davis, 1989).

Figure 3: User acceptance of information technology

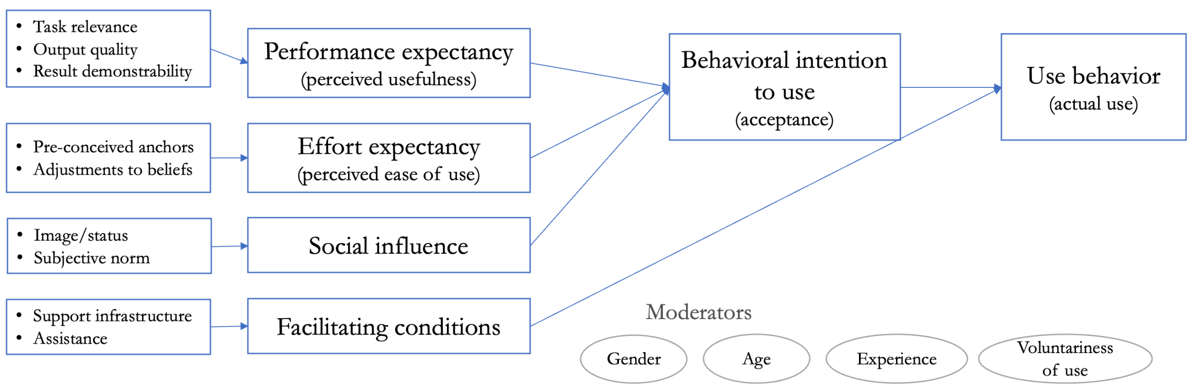

Source: Venkatesh *et al* (2003).

When it comes to the actual use of the algorithm, two interviewees did not use the algorithm at all, while two others indicated that they used it in about 15 percent of the cases they handled. Estimated time savings reported by these two interviewees were very minimal, about 10 minutes per usage. For an average of four cases a week out of which at most one is sped up by the use of the algorithm, this means a maximum of 10 minutes saved per week. This seems to contrast with the bigger reported

[6] For an application of this model to the acceptance of AI in the workplace, see Hoffman and Nurski (2021b).

productivity increases mentioned earlier in this report. However, it is important to note is that all respondents were fairly experienced workers, who had less benefit from using the technology (see below). All of the interviewees agreed that usage and related time savings are greater for less-experienced workers, such as interns or new colleagues, and for cases outside one's usual area of expertise, for example when covering for an absent colleague with a different specialisation.

3.3 Barriers to the use of the algorithm

Early studies showed that usefulness is a stronger predictor of uptake than ease-of-use: *"Users are often willing to cope with some difficulty of use in a system that provides critically needed functionality"* (David, 1989). In the current case, we identified two main barriers, one in the ease-of-use and another one in the usefulness of the AI algorithm. A minor barrier was the lack of social norm for using the AI. Table 3 summarises the drivers and barriers identified in the case of FIT, and groups them according to the framework of Venkatesh *et al* (2003), illustrated in Figure 3.

Table 3: Identified drivers and barriers to worker acceptance of the AI product at FIT

	Identified drivers	Identified barriers
Ease-of-use	AI-retrieved answers are easily located	AI-retrieved answers cannot easily be altered (PDF)
Usefulness	Task takes up large share of workload	AI-retrieval works best in cases where also the worker's memory works best
Social influence	Some positive self-image associated with use	No social norm for using the AI
Facilitating conditions	Workers know where to get assistance	

Source: Bruegel based on interviews at FIT [12] and Venkatesh *et al* (2003).

3.3.1 Perceived ease-of-use: a technology design issue

The AI developer explicitly aimed to focus on the user experience from the start by having a working solution, integrated into the existing IT systems early on in the development process. The AI-suggested answers were integrated in the CRM system (Microsoft Dynamics) as an extra tab in the screen displaying the information on the incoming question. Interviewees indicated that they can find the suggested answers easily:

Interviewee 1:

"You can find [the AI suggestions] very quickly. I go to the question and then I have a separate tab [in the CRM] where I immediately see the answer options. So with one click I can open them. [...] You can actually see fairly quickly from the question whether the answer is going to fit or not."

Interviewee 2:

"It's literally pushing a button, it can't be any easier than that. So I think in terms of ease of use it's really top notch."

Interviewees indicated that it is easy to copy text from an old answer to a new e-mail, even though the old answer usually needs some editing to tailor it to the new client's question.

Interviewee 1:

"So I find that in most cases you can't literally copy paste [the AI-retrieved past answer], even with questions that are asked often. [...] You always have to tweak that answer a little bit. So I copy the text and then I adjust it as I want for the new client. It's not that I just copy the whole answer and then send it to the company."

The retrieved past answers shown by the AI include both the text that was sent in an email to that previous client and any PDF attachments that were included with that email and other information on the previously answered case (such as name, sector and country of the old case). These PDF attachments usually contain lists of contact details or event dates. While the text can be easily copied from an old answer to a new email draft, the PDF cannot be edited in the CRM system itself. For that, the case handler needs to find the original file (Word or Excel) from which the PDF was generated. In the current workflow, most case handlers store these original files in the collective document storage system (Microsoft SharePoint). So to edit these files, they need to retrieve the right folder on the SharePoint, find the original document from which the PDF was generated and then update the document and regenerate the PDF.

Interviewee 1:

"I select the text of the answer if I want to copy it. But the address list, for example [this is an example of a PDF attachment, ed.], I do download that, so I can adjust it if necessary. Because I can't make adjustments [to the PDF attachment] in the [CRM] system if the answer has already been sent.”

Interviewee 3:

[The AI-retrieved answers are] in the overview, so that’s easy to find. […] But you then have to click on that link [to the old answer] and then you go to that case, still in Dynamics [the CRM system]. […] Then you have the answer as an attachment. […] And then you have to download that attachment and normally that’s a PDF. So not a file that you can edit. So there are I think 2-3-4 steps. Whereas, according to my usual way of working in Sharepoint, it's faster for me to look at the AI suggestions in the CRM and then just search for the old case directly on Sharepoint.

Interviewee 4:

It is in itself very user-friendly to access the AI suggestions. It’s easy to find, absolutely. The problem is you often still have to modify the information in attachment, so people think "then I better start from scratch or do it on the SharePoint" because it’s all automatically stored there.

Not being able to edit attachments to previous answers in the AI-suggested list was therefore a main barrier to the ease-of-use of the AI system. The AI project lead had anticipated this risk and wanted to mitigate it by getting users to store the attachments directly in the CRM instead of on the SharePoint. However, end users did not change their habits or workflows because they still very much preferred the ease-of-use of the SharePoint. When ease-of-use is insufficient, end users will use the software differently than the designers intended.

Interview AI lead:

“We actually want to discourage employees from storing information separately on the SharePoint. In principle, case handlers should store all case information in the CRM. But people don’t do this enough because SharePoint is much easier to use. So that's also something, how can we make sure it's easier to use the CRM.”

3.3.2 Perceived usefulness: an organisational or job design issue

All four interviewees agreed that answering trade questions takes up a large part of their workload. This activity on average takes up 60 percent to 70 percent of their time, but can vary during the year depending on the presence of other time-consuming activities such as hosting events or organising trade missions. The large portion of time spent on answering trade questions indicates that there was indeed a significant potential for time savings and productivity increases in this area, making the AI algorithm very relevant for their tasks.

The main hindering factor limiting the usefulness of the AI-supported retrieval of previously answered questions was the types of cases for which the AI worked best. The AI works best when it has seen several examples of similar questions before, ie the more 'repetitive' type of questions. In those cases, the AI can retrieve reusable answers from the recent past. However, in those exact same cases, the more repetitive ones, also the experienced worker himself can easily remember that they answered similar questions in the past and can find the information that was used in those past answers on the SharePoint:

> Interviewee 3:
>
> "We get several questions around the [redacted] industry every year, which means I automatically know where to find the information about this industry. So I know a similar question came in 1-2 months ago. Then I just go directly to the SharePoint folder where that other case is, without checking the list of [AI] inputs."

A worker called this their "historical memory". Several interviewees indicated that newer, less-experienced colleagues who didn't have a long history of answering questions, would be able to benefit from the AI-assisted retrieval more.

> Interviewee 3:
>
> "I explained [to my new colleague] where the cases are and where our older answers to previous cases are. But as you can imagine she has no historical memory of "ok, I remember in 2019 we got a similar question like this". So I suspect that for her such the AI suggested answers may be even more important, even for questions that may be so simple for me that I don't need to check the AI."
>
> Interviewee 1:
>
> "But just for finding things faster or indeed for interns or new colleagues, that's where I see the added value of it especially."

Finally, one interviewee distinguished between three types of questions: (1) those that are so frequent (very routine questions) that they just handled them last week, so they don't need the assistance of the algorithm; (2) those that are so rare (unique questions) that the algorithm can't help, because no similar question has been answered in the past; (3) those in the middle (somewhat routine questions) that may have been answered at some point in the past, but the worker doesn't immediately remember.

> Interviewee 4:
>
> "If I don't remember where to look in the Sharepoint, but I do think we must have had a similar question like that before, then I'm going to use the algorithm especially. It could be that I just don't find the old case by myself, or that a colleague has saved the information about this case on their desktop. In that case, I cannot find the old information on the SharePoint, but the AI can find it in the CRM."

Table 4 summarises the perceived usefulness of the AI assistance in relation to the task routines on the one hand and the experience of the worker on the other hand.

Table 4: Usefulness of AI support by 'task routineness' and 'worker experience'

		Task routineness		
		Very routine question	Somewhat routine question	Unique questions
Frequency		Every week/month	Every year	Once
Quality of AI retrieval		Highest	High	Low
Usefulness of AI to worker	Low experience worker	Highest	High	Low
	High experience worker	Low (worker remembers by themselves)	Highest	Low (no past answer exists)

Source: Bruegel based on interviews [12].

3.3.3 Social influence and facilitating conditions

All interviewees agreed that there was no social pressure or expectation from colleagues or managers to use the AI tool, nor were there any compulsory policies to use these tools. In fact, interviewees considered most of their colleagues or managers to be less interested in, or less capable of, using new technologies compared to themselves.

Interviewee 3:

> "My colleague is also kind of old school. Technologically equipped, but not really fanatical. We also have no real official policies that make the tool compulsorily to use. It was just offered as a nice additional help."

Interviewee 4:

> "My manager knows that I will spontaneously use the AI when it is necessary for my work. I don't feel they have to direct me in that."

Interviewee 2:

> "[My manager] and I have a very different perspective on new applications. I know [they are] often very sceptical. So I have a gut feeling that [they] would not be very enthusiastic about it."

Interviewee 2:

> "I know that my manager certainly found the AI interesting, [they are] really very open to new technologies. But, [...] if the head of the team doesn't use it, it's not going to be implemented among assistants."

Some interviewees appropriated a positive self-image from learning to use frontier technologies, contrasting themselves with others who they consider less-proactive personality types.

Interviewee 4:

> "I see that – in an ideal world – the algorithm could ensure that you have more time for research and you can learn more. But that effect will be different according to people's personality types. Because, whichever way you turn it, there will still be people who just do their job as it should be done, but who are not going the extra mile and do not proactively learn or search."

All interviewees understood clearly where they could get assistance and support for using the algorithm. When asked where they would direct questions about the AI system, all of them mentioned their general IT support or the AI project lead or both. One person also mentioned the information session they had followed on the AI tool. All were confident they would be helped if they encountered any issues, but none had requested any help.

Interviewee 2:

"I actually knew perfectly who could help me. It's intuitive. I would send a message to [the AI lead] or to [the IT support …] So I believe that enough people in FIT knew enough about that tool to guide me."

4 Impact and support

4.1 Impact on work divisions, learning and social relationships

Given that the usefulness of the AI support varied by task routineness and by worker experience (see 3.3.2), some managers or senior team members distributed the more routine questions – where the AI-retrieval works best – to less experienced colleagues.

Interviewee 3:

"I think maybe unconsciously I take the more substantive questions for myself and then pass on the more list-like questions to my colleague or to interns."

If this practice systematically took place, it could undermine the learning and development opportunities for junior or less-experienced workers in the long run. Other managers therefore intentionally didn't pass on the most routine – AI-assisted – questions to their junior team members, and they specifically mentioned their intentions with respect to learning. They also made sure to provide enough variety in the questions assigned to less-experienced workers for the same reason.

Interviewee 4:

"Of course, it depends on their capacity and skills, but I do try to push my junior colleagues out of their comfort zone. I also try to give them more difficult cases and then follow-up as co-handler. […] I want them to learn in their work. So I do try to see if they've had a lot of similar questions recently. Then I will do that case myself and give them something else. […] So it's not necessarily the cases that could be answered by the algorithm that I hand off to the assistants. Absolutely not."

The question here becomes what the 'learning' of the workers using the AI needs to be: do they need to learn the skills required to use and check the AI, or do they still need to learn the underlying domain-specific knowledge as well? In this specific case study, workers need to be able to check the quality of the AI suggestions as well as to manually complement any missing information for very unique questions that the AI cannot answer. They should therefore still have sufficient domain expertise in order to understand and complement the limitations of the algorithm.

The increased digitalisation of the work meant that the CRM and the AI became a sort of Knowledge Management System (KMS) that stored the collective knowledge of the workers in a digital system. According to the HR lead, this led to a decrease in within-team communication and collaboration.

> HR lead:
>
> "[I asked the case handlers:] 'You probably hear each other quite a lot, right? Because those AI-suggestions have to be tailored to the new case.' Their answer was 'yes but we talk much less than before. We talk via the system.' What they actually wanted to say was, on a collaborative level and on a mental level, we have gone backwards. We have progressed in efficiency to improve and facilitate our work. [...] but we are constantly in that tool and we speak through the tool."

Although this was not true for everyone. Some workers still very much preferred to share knowledge about cases in person.

> Interviewee 4:
>
> "If I think a colleague has previously answered a similar question and we are in the office together at that time, my first instinct would still be "hey did you answer that case?". That's always going to be my first instinct [before checking the AI].

On a side note, the same employee also mentioned that the impact of remote work on collaboration was bigger than the impact of the KMS.

> Interviewee 4:
>
> "Of course, if you work from home it's slightly different. Then you don't send a message on Teams right away and you look for the answer by yourself first."

4.2 Path towards increasing AI acceptance

Involvement of HR in the AI adoption strategy focused on the creation of new roles and functions. This started with the instatement of the function 'project manager data architecture & artificial intelligence' (referred to in this report as 'AI lead') in 2017. To support the transformation to the data-driven culture, voluntary roles of 'data stewards' were created throughout the organisation, at this point without any links to formal HR processes such as renumeration or performance evaluation. Three people volunteered for these roles. At the same time, a different process was running in which a strategic personnel plan was developed. This plan only had budget for two promotion functions with pay grades and function descriptions for these data stewards. When the third phase of AI adoption (see 2.1.4) led to the development of a better data infrastructure, it was also decided in the strategic personnel plan to create and hire for a new role of Chief Data Officer.

> HR lead:
>
> "Three people in our organisation have taken on the role of data steward, out of enthusiasm, out of wanting to contribute to the future. In the strategic staffing plan it later turned out that there were only two promotion opportunities, so then... two does not equal three."

The HR lead further detailed the two different speeds at which formal strategic HR exercises run, compared to the fast progress that data-driven and AI discovery processes tend to make.

> HR lead:
>
> "But you have different speeds at work in the organisation and a strategic [personnel] plan has a certain duration because that is a very bottom-up exercise. Because of its big impact, it also has to be well thought through. Then you have the progression of projects [that move faster, ed.]. That's a natural thing in an organisation. You want your projects to move ahead, and the anchors afterwards [ie HR processes] are at a different level or at a different stage so it doesn't track together."

To support the worker in their adoption and acceptance of AI, the HR and AI project leads focussed on communication, organisational development and individual learning and development.

Communication had been a key focus point from the start. The goal was to address the perceived ease-of-use driver (in the technology acceptance framework of 3.2) and general fear surrounding the use of AI tools.

> Interview AI lead:
>
> When it comes to AI, I may be exaggerating, but we have spent at least as much time on awareness as on technical development itself, precisely to maximise acceptance and to really

show what AI is. I even remember a meeting where a colleague burst into tears because she read so much in the media about AI and was really frightened by it. Then we also gave a presentation about the media's perception of AI compared to the reality and where we are going. [...] That did do a lot of good.

The IT department also has regular meetings with the business departments which also focussed a lot on joint communication efforts.

Interview AI lead:

We have a business-IT department meeting, which includes a number of IT representatives and a number of people delegated from the business department. A lot of emphasis is put on communication in this meeting. How are we going to communicate this? There is now a new change, how should we communicate it? [...] We find that it is super important. We can develop the best applications, but if we don't communicate well, they won't be used and it's almost lost money.

On the side of organisational development, parts of the central IT support function were decentralised, ie some tech support tasks were moved from a central organisational entity to the frontline teams. This included some IT support to FIT's own employees (employee tech support) in the form of new roles of 'key users' that could act as first contact points the use of internal digital tools including the CRM and the AI. The decentralisation also included the IT support provided to FIT's external clients when using digital tools such as question forms or event registrations (customer tech support). This decision clearly made the support infrastructure and assistance more accessible (ie addressing the 'facilitating conditions' driver in the technology acceptance framework in section 3.2).

Interview HR:

"We have two positions now in which people are going to be facilitators for teaching those tools. They will take on a senior role in that function and in that senior role they will also be expected to ensure that the team is involved in the usage of tools. This way, adoption by the users can take place in the team itself, so not from a support service or from the project, but from the group itself."

Interview HR:

"Then on the organisational development side, we are setting up a business IT helpdesk so that the focus of IT and the workload there can also be shifted slightly to the front line [ie the case handlers]. Certain questions from companies that are now handled by IT can then be taught in the front line."

Finally on the individual learning and development side, beyond the training plans that were designed for each of the different groups of employees (see section 2.2.3), HR also set up a reward system that rewards employees in non-monetary terms for achieving milestones in digital tool usage. Such rewards could for example include a solar-powered power bank for charging mobile phones and laptops. This reward system aimed to align incentives between the employer and its employees; award social status to digital tool usage; and install a social norm for tool usage that had been missing in the organisation so far (ie addressing 'social influence' driver in in the technology acceptance framework in section 3.2).

Interview HR:

"This reward system is put into the learning path, every time a milestone is achieved in learning the tool, to motivate them. If they have X number of points, they can exchange that for a gift."

The final step they arrived at – at the moment when this case study was conducted – was to really focus on addressing (the perception of) the usefulness of the AI or other digital support tools, in terms of task relevance, output quality and results demonstrability (in in the technology acceptance framework of 3.2).

Interview HR:

"I am now on the path of Jane Hart [holistic Learning & Development guru, ed.]. How can we get employees to use our tools? [...] Basically, the bottom line is that we really need to go to impact and start from there to effectively move people towards acceptance."

5 Conclusion and recommendations

This case study shows how AI support in frontline processes does not necessarily have to displace or control workers. Some crucial factors in successfully adopting a human-centred approach to AI were identified. First, a 'fast-discovery' phase – including the development of fully operational proofs-of-concept – can facilitate the trialability of AI, ie the ability to experiment with an innovation before fully committing to it. This in turn reduces uncertainty and facilitates adoption. Second, workers' assessments of the potential productivity gains and the current data availability are crucial for assessing both the technical complexity and business value of the opportunities, even before any development takes place. Involvement of workers in the evaluation and improvement of the early concepts also significantly increase the quality of the algorithm after the first minimum viable product are released. Third, it is important to involve human resources managers in the process of technology adoption early on, to facilitate the alignment between HR processes, IT processes and business processes from the start. Finally, financial support from diverse European and national subsidy channels (specifically the EU's SRSP and ERDF programmes) were specifically targeted to assist in the adoption of AI.

As this case study has also demonstrated, making AI support available to workers is not sufficient to ensure its widespread usage throughout the organisation. More experienced workers in particular were not very inclined to make much use of the algorithm, mostly because they had less need for the AI support. The technology needs to be seamlessly integrated into workflows to persuade workers to optimally use it. This is a technology design issue that is well understood – also well recognised by the participants in this case – and can be easily addressed. A bigger barrier to worker acceptance is the usefulness of a new technology for a specific task to a worker. The potential of AI to support work is usually framed in relation to the routineness of tasks, but as this case study has shown, it also needs to be assessed in the context of worker experience and task allocations (or work divisions) among workers. Depending on the interplay between task routineness and worker experience, the new technology might necessitate a rethinking of the work division between technology and workers on the other hand, and junior and senior workers on the other hand. Contrary to the beforementioned technology design issue, this is an organisational design issue.

Successful human-centred AI adoption will therefore depend on the availability of investments in complementary intangible organisational capital. These investments go beyond training in digital skills. They are investments into the redesign of organisational processes that are necessary to reap the benefits of new technologies. This includes both individual people-management processes, like hiring, performance evaluations and reward systems, and organisational development processes including business process reengineering and organisational redesign. These types of investments are currently poorly captured in company balance sheets (Brynjolfsson *et al*, 2021) and as such, we have poor information about intangibles in national accounts as well. In order to stimulate research and understanding in this area, it would be advisable to adapt accounting standards to capture firms' intangible investments in their people and organisations.

References

Alderucci, D., L. Branstetter, E. Hovy, A. Runge and N. Zolas (2020) 'Quantifying the impact of AI on productivity and labor demand: Evidence from US census microdata', mimeo, ASSA 2020 Annual Meeting, Allied Social Science Associations

Atkin, D., A. Chaudhry, S. Chaudry, A.K. Khandelwal and F. Verhoogen (2017) 'Organizational barriers to technology adoption: Evidence from soccer-ball producers in Pakistan', *The Quarterly Journal of Economics*, 132(3): 1101-1164

Baker, J. (2012) 'The Technology–Organization–Environment Framework', in Y.K. Dwivedi *et al* (eds) *Information Systems Theory: Explaining and Predicting Our Digital Society, Vol. 1*, Integrated Series in Information Systems, New York, NY: Springer

Baker, G.P. and R. Gil (2012) 'Clinical papers in organizational economics', in R.Gibbons and J. Roberts (eds) *The handbook of organizational economics*, Princeton University Press

Bloom, N., R. Sadun and J.V. Reenen (2012) 'Americans do IT better: US multinationals and the productivity miracle', *American Economic Review*, 102(1): 167-201

Bresnahan, T.F., E. Brynjolfsson and L.M. Hitt (2002) 'Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence', *The Quarterly Journal of Economics*, 117(1): 339-376

Brynjolfsson, E., D. Rock and C. Syverson (2021) 'The productivity J-curve: How intangibles complement general purpose technologies', *American Economic Journal: Macroeconomics*, 13(1): 333-72

Damioli, G., V. Van Roy and D. Vertesy (2021) 'The impact of artificial intelligence on labor productivity', *Eurasian Bus Rev* 11: 1–25, available at https://doi.org/10.1007/s40821-020-00172-8

Davis, F.D. (1989) 'Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology', *MIS Quarterly*, 13(3): 319–340, available at https://doi.org/10.2307/249008

Devaraj, S. and R. Kohli (2003) 'Performance impacts of information technology: Is actual usage the missing link?' *Management Science*, 49(3): 273-289

European Commission (2020) *European Enterprise Survey on the Use of Technologies Based on Artificial Intelligence: Final Report*. European Commission DG Connect, available at https://data.europa.eu/doi/10.2759/759368

Giovannucci, D. (ed) (2004) *National trade promotion organizations: their role and functions*, World Bank Group

Hoffmann, M. and L. Nurski (2021a) 'What is holding back artificial intelligence adoption in Europe?' *Policy Contribution* 24/2021, Bruegel

Hoffmann, M. and L. Nurski 'Workers can unlock the artificial intelligence revolution', *Bruegel Blog*, 30 June 2021, available at https://www.bruegel.org/blog-post/workers-can-unlock-artificial-intelligence-revolution

Jaramillo, C. (1992) 'The basic functions of national trade promotion organizations', *International Trade Forum* 3(Jul-Sep): 18

Jauk, S., D. Kramer, A. Avian, A. Berghold, W. Leodolter and S. Schulz (2021) 'Technology Acceptance of a Machine Learning Algorithm Predicting Delirium in a Clinical Setting: a Mixed-Methods Study', *Journal of Medical Systems* 1;45(4):48, doi: 10.1007/s10916-021-01727-6

Kawaguchi, K. (2020) 'When Will Workers Follow an Algorithm? A Field Experiment with a Retail Business', *Management Science* 67(3): 1670-1695

Landauer, T.K. (1995) *The trouble with computers: Usefulness, usability, and productivity*, MIT Press

Lundblad, J.P. (2003) 'A Review and Critique of Rogers' Diffusion of Innovation Theory as It Applies to Organizations', *Organization Development Journal*, 21(4): 50

Rogers, E.M. (1983) *Diffusion of Innovations*, New York: The Free Press, Macmillan Publishing

Rueda-Cantuche, J.M., P. Piñero and Z. Kutlina-Dimitrova (2021) *EU exports to the world: effects on employment*, EUR 30875 EN, Publications Office of the European Union, Luxembourg, doi:10.2760/556206, JRC126534

Venkatesh, V., M. Morris, G. Davis and F. Davis (2003) 'User Acceptance of Information Technology: Toward a Unified View', *MIS Quarterly*, 27(3): 425-478

WTO Stats dashboard, World Trade Organisation, available at https://stats.wto.org

Xu, Y. and L. Zhu (2021) 'Technology Adoption: The Impact of Employee Incentives', mimeo, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3453953

Zanatta, M., I. Costa and S. Filippov (2006) 'Foreign direct investment: Key issues for promotion agencies', *Policy Brief* 10, United Nations University

Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. Beede, C. Buffington, N. Goldschlag, L. Foster and E. Dinlersoz (2020) 'Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey', *NBER Working Paper* No. 28290, National Bureau of Economic Research

Annex: List of case study materials

[1] https://corporate.flandersinvestmentandtrade.com/en/results retrieved on July 11th 2022

[2] https://corporate.flandersinvestmentandtrade.com/en/news/2021-was-record-year-flanders-exports retrieved on July 11th 2022

[3] https://corporate.flandersinvestmentandtrade.com/en/strategy retrieved on July 11th 2022

[4] Slide deck 'In-depth data & AI' by FIT & interview with AI lead on 4/11/2021

[5] https://overheid.vlaanderen.be/nieuws/vlaamse-regering-keurt-vlaamse-datastrategie-goed

[6] https://www.vlaanderen.be/digitaal-vlaanderen/vlaamse-digitale-strategie/desi

[7] Slide deck 'AI to improve workflow at FIT' by Radix

[8] https://radix.ai/cases/how-ai-assisted-question-answering-helps-flanders-attract-foreign-investors/

[9] https://ec.europa.eu/reform-support/blue-print-artificial-intelligence-ai-adoption-flanders-investment-trade_en

[10] Interview with FIT AI lead on 23/11/2021

[11] https://corporate.flandersinvestmentandtrade.com/nl/nieuws/hoe-fit-coronatijden-handelspromotie-een-boost-geeft-met-artificiele-intelligentie-ai

[12] Interviews with four case handlers in May and July 2022

[13] Interview with HR lead on 18/05/2022

[14] https://radix.ai/discover

bruegel

Bruegel, Rue de la Charité 33, B-1210 Brussels
(+32) 2 227 4210
info@bruegel.org
www.bruegel.org