Technology Boom, Labor Reallocation, and Human Capital Depreciation*

Johan Hombert† Adrien Matray‡

Abstract

.

*We thank and Pierre Azoulay, Laurent Bach, Nick Bloom, Denis Gromb, Nandini Gupta, Gregor Jarosch, Stefan Lewellen, Stefano Lovo, Marco Pagano, Richard Rogerson, Chenzi Xu as well as seminar participants at HEC-Paris, BI-Oslo, ESSEC, NOVA Business School, IFN Stockholm and conference participants at the 14th Csef-Igier Symposium on Economics and Institutions, the American Finance Association 2019, AFFI 2018, ESSFM (Gerzensee) Symposium 2019, Labor and Finance Conference (Chicago), the LBS Summer Symposium. Hombert acknowledges financial support from the Investissements d'Avenir Labex (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047). Matray acknowledges financial support from the Griswold Center for Public Policies and the Louis A. Simpson Center for the Study of Macroeconomics.

†HEC Paris and CEPR, hombert@hec.fr

‡Princeton University and CEPR, amatray@princeton.edu

1 Introduction

Technological change over the past thirty years has been marked by the development of Information and Communications Technologies (ICTs) that fostered automation and digitization. These technologies benefited high-skill workers and accelerated the fall in the employment in middle-wage, routine occupations, promoting the concept of skill-biased technological change and job polarization (e.g., David, Katz, and Kearney, 2006; Goos, Manning, and Salomons, 2009; Jaimovich and Siu, 2020; Jaimovich, Saporta-Eksten, Siu, and Yedid-Levi, 2021). But is technological change always beneficial to high-skill workers, in particular when their human capital is tightly linked to the development of the new technologies? Before ICTs pervade all sectors of the economy, their fast-paced development in the 1990s led to high labor demand and high wages in the ICT sector, making careers in the ICT sector appear attractive at the time and indeed attracting many high-skill workers. For instance, a third of the cohort of high-skill individuals who entered the labor market in the late 1990s started their career in the ICT sector in France (see Section 2). The same pattern holds in the US (see Appendix C).

Despite the quantitative importance of this reallocation of high-skill labor, little is known about its long-term consequences for the large number of high-skill workers who went to the booming ICT sector and acquired skills tied to the new technologies they contributed to develop. Once workers invested in the required skills, what happened to their human capital? Did fast-paced technological change during this period make these skills quickly obsolete, or instead, did workers who acquired early-on skills linked to the new technologies increase the long-term value of their human capital?

Contributing to develop new ICTs has an a priori ambiguous effect on workers' long-term human capital. If it allows workers to acquire useful skills, their human capital will benefit in the long run, leading to higher wages. This would happen if these skills are firm specific and make workers more productive conditionally on staying at the same firm; or if the skills are task specific and transferable to other firms or other occupations, insulating workers from firm idiosyncratic shocks (e.g., Gathmann and Schonberg, 2010). If instead, their human capital depreciates quickly because of fast-paced technological change during the ICT boom, they will end up with lower levels of human capital and lower wages. Such obsolescence would imply that skills are not only *task* specific but also *vintage* specific (Chari and Hopenhayn, 1991; Deming and Noray, 2020).

To shed light on this question, we study the potential human capital depreciation of high-skill workers who start their career in the late 1990s booming ICT sector, by analyzing their wage dynamics. In particular, we study how it compares to the wage dynamics of workers also performing high-skill technical tasks, but in other sectors. And how it compares to the wage dynamics of workers also going to the ICT sector but after the boom, when the technology has stabilized and new vintages of technology appear more slowly.

We implement our tests by using French administrative matched employer-employee data for the period 1994–2015, which we link to the universe of firms' financial statements from tax filings. The data contain high quality, longitudinal information on workers' wages and career paths and firms' characteristics, which allow us to partition workers based on the characteristics of the firm they work in, and the qualification and occupation they have. We use the combination of the worker's occupation and the employer's characteristics to infer the tasks workers perform and the extent to which these tasks are related to new ICTs. The assumption behind our empirical strategy is that the skills developed by a STEM worker at (for instance) Microsoft during the late 1990s are more tightly linked to ICT than those of a STEM worker at Airbus.

A preliminary contribution of the paper is to analyze the margins of skilled labor reallocation during a technology boom. We find that almost all the reallocation takes place at the extensive margin of the labor market. The share of high-skill labor market entrants starting in the ICT sector almost doubles during the boom, from 17% to 31%, before dropping back to 19% when the boom ends. By contrast, the net flow of seasoned workers barely contributes to high-skill employment growth in the ICT sector during the boom. This result has two implications. First, conceptually, it points to the existence of different vintages of human capital within occupations. Workers with older vintages of human capital outside the ICT sector appear to be poor substitutes for new workers in the ICT sector during the development phase of new technologies, even if both old and new workers have the same occupation recorded in administrative data (e.g., IT engineers). Second, empirically, the sharp delineation in time of the period of intense experimentation in the ICT sector helps for identification. It allows us to identify a priori which cohorts have their human capital tightly linked to the development of new technologies, which cohorts were already in ICT before these technologies evolved rapidly, and which cohorts joined ICT after technologies stabilized.

We then turn to our main question: how does post-schooling human capital accumulate and depreciate as a function of the technological content of the tasks workers have to perform and the vintage of entry in the labor market?

To guide the empirical analysis, we outline a simple dynamic two-sector model with overlapping cohorts, worker sectoral choice and on-the-job human capital accumulation/depreciation. The model provides an intuitive decomposition of the average wage in a sector-cohort into three components: the (sector-specific) wage rate that reflects labor supply and demand in the sector, the (sector-cohort-specific) human capital accumulated since labor market entry, and a (sector-cohort-specific) selection term that depends on the endogenous sorting of workers across sectors. The model shows that the human capital accumulation component can be backed out by comparing the wage dynamics across cohorts of workers and across sectors, i.e., by running difference-in-differences regressions.

In the estimation, the first difference compares in any given year workers who started in the ICT sector to similar workers from the same cohort, who started in other sectors. This allows us to include cohort×year fixed effects and ensures that we compare workers exposed to the same macroeconomics shocks at the same stage of their career. In particular, cohort×year fixed effects absorb the well-documented impact of macroeconomic conditions at the time of labor market entry on long-term earnings. The second difference compares in any given year workers who started in the same sector, but across cohorts. This allows us to include starting sector×year fixed effects and ensure that we compare workers exposed to the same sector-specific shocks.

Our central result is that workers who enter in the ICT sector during the boom earn 5% higher entry wages on average but end up with long-term wages 6% lower 15 years out, relative to workers from the same cohort with similar characteristics starting in other sectors. This result holds both in a simple-difference specification across sectors, which compares workers of the boom cohort starting in ICT versus outside ICT, and in the difference-in-differences specification across both sectors and cohorts, which additionally uses workers of the post-boom cohort as a control group. The point estimates in the simple difference and in the difference-in-differences are similar. It implies that the effect comes from the boom cohort of ICT workers experiencing a long-term wage decline, while the wage dynamics of the post-boom cohort of ICT workers does not stand out as particularly positive or negative. Therefore, the result is not explained by a persistent imbalance between labor supply and demand in the ICT sector or a permanent shift in sector productivity. Instead, low long-term earnings concentrated on the boom cohort of ICT workers is consistent with steady depreciation of human capital accumulated in the ICT sector during the boom.[1]

1. The pattern is similar for cumulative wages, implying that reverse backloading, whereby

A remaining confounding factor is endogenous sorting of workers with different ex-ante productivity across sectors, places, firms and occupations. We progressively saturate the specification with an increasing set of fixed effects, which relax the identifying assumption under which changes in long-term wages identify the effect of human capital depreciation. Workers' ex-ante productivity could be correlated with the sector of entry, which we control for by including starting sector×year fixed effects. Workers might also sort across places based on their ex-ante productivity, which we address by adding local labor market×cohort×year fixed effects, or across occupations, which we control for with occupation×cohort×year fixed effects. High-productivity workers might also sort into high-productivity firms and this sorting behavior might vary across cohorts. We address this potential problem by including worker fixed effects and controlling for multiple firm characteristics interacted with cohort×year fixed effects. High-productivity workers may also be on different earnings paths. We control for this by including quintile of entry wage×year fixed effects.

These specifications control for most sources of biases that might be introduced by sorting into jobs correlated with unobserved differences in workers' wage level. Our model helps clarifying that the only source of remaining bias is that unobserved worker heterogeneity shifts not only the wage *level*, but also wage *growth*.[2] We address this problem by bringing a new cohort to the analysis: the cohort of workers who sort into jobs just before the ICT boom started. Under the assumption that these workers did not sort into ICT based on the anticipation of the subsequent boom, these workers constitute a group whose sorting decisions could not have been affected by the boom. Using the same difference-in-differences framework as before and comparing the pre-boom cohort with the post-boom cohort, we show that workers of the pre-boom cohort experience quantitatively the same long-run wage discount as workers of the boom cohort. The fact that the pattern of wages of the pre-boom cohort cannot be explained by sorting based on the presence of the boom (by construction) reinforces the interpretation that the long-run wage discount is due to human capital depreciation.

Finally, we study potential mechanisms by which human capital accumulated during the boom depreciates at an accelerated pace. First, technological change

workers are paid below their productivity in the future in exchange for higher pay early in their career (Lazear, 1981) does not explain the results.

2. This distinction is important because the literature on work conditions and workers' wage dynamics, such as the literature on peer effects, has focused on selection issues coming from unobserved differences in wage level but has overlooked the issue of selection coming from unobserved differences in wage growth.

may accelerate during the boom, resulting in faster obsolescence of skills tied to the evolving technology (Chari and Hopenhayn, 1991; Deming and Noray, 2020). Second, workers losing their jobs in the bust may lose firm-specific human capital or be poorly matched later on and end up on a different career path associated with lower long-term earnings (e.g., Gibbons and Katz, 1991; Jacobson, LaLonde, and Sullivan, 1993; Wachter and Bender, 2006; Burdett, Carrillo-Tudela, and Coles, 2020).

We find support for the skill obsolescence mechanism: the wage discount is larger for workers whose jobs involve performing tasks that are more tightly connected to the development of new ICTs. We establish this result in two ways. First, within similar broad tasks (STEM versus non-STEM occupations) but across skill levels, we find that the wage discount is concentrated on high-skill workers, while by contrast, lower-skill workers in ICT firms have similar wage dynamics as lower-skill workers in non-ICT firms. This is consistent with the fact that lower-skill workers accomplish more routine tasks, which are less dependent on the frontier technologies currently being developed. Second, within high-skill workers, we find that the long-run wage discount is larger in firms and sectors that are more technology intensive, as measured by the share of high-skill workers in their workforce.

By contrast, we find no support for the job termination mechanism. We first show that while workers who start in ICT during the boom are more likely to experience a forced job termination, the effect is several orders of magnitude too small to explain the long-term wage discount. We also show that controlling for job termination explains a negligible part of the lower wage growth, implying that the long-term wage discount is not explained by workers losing firm-specific, non-portable skills.

Related literature. Our first contribution is to the literature on skilled-biased technological change. This literature has shown how technological change has affected the usage of tasks (e.g., Autor, Levy, and Murnane, 2003; Spitz-Oener, 2006; Goos, Manning, and Salomons, 2014; Kahn and Hershbein, 2018; Blair and Deming, 2020; Jaimovich and Siu, 2020; Jaimovich, Saporta-Eksten, Siu, and Yedid-Levi, 2021) and the value of human capital (e.g., Beaudry, Doms, and Lewis, 2010; Beaudry, Green, and Sand, 2016). A common theme in this literature is that high-skill workers are the winners of technological change. We add to this literature by showing that, within high-skill workers, technological change has a negative impact on workers who contributed to its development and diffusion.

Second, we contribute to the small and nascent literature studying how firm-level innovation gains are shared with workers (Aghion, Akcigit, Hyytinen, and Toivanen, 2018; Kline, Petkova, Williams, and Zidar, 2019; Aghion et al., 2020) and more specifically with innovators (Toivanen and Väänänen, 2012; Depalo and Di Addario, 2014; Bell, Chetty, Jaravel, Petkova, and Van Reenen, 2018; Kogan, Papanikolaou, Schmidt, and Song, 2019). This literature analyzes the impact of differences in innovation across firms at a given point in time (i.e., for a given vintage of technology) on workers' earnings. We add to this literature by showing that a wave of innovation has a negative impact on the earnings of skilled workers who contribute to its development, because their vintage of human capital becomes obsolete.

Our third contribution is to the literature studying the extent to which skills are specific to, or portable across different uses, such as across occupations (Gibbons, Katz, Lemieux, and Parent, 2005; Kambourov and Manovskii, 2009), sectors (Dustmann and Meghir, 2005) and firms (Card, Devicienti, and Maida, 2014; Card, Cardoso, and Kline, 2015; Arellano-Bover and Saltiel, 2020). Our contribution to this literature is to show that skills are also specific to the *vintage* of technology they rely on, and that this dimension of skill specificity is quantitatively important. The specific implication of vintage-specific skills is that accelerating technological change depreciates workers' human capital tied to technology, which affects wages and can weight on workers' job mobility.

For this reason, we also contribute to the literature on vintage human capital, which proposes that several vintages of knowledge can co-exist, and that technological change makes old vintages obsolete (Chari and Hopenhayn, 1991; MacDonald and Weisbach, 2004). Deming and Noray (2020) provide evidence from job vacancy data that skill requirements of STEM occupations can change rapidly, making seasoned workers' skills obsolete. We show that skill obsolescence is stronger for the large cohort of workers joining the booming technology sector. Thus, technology booms triggering temporary changes in labor sectoral allocation can have long-lasting consequences through the human capital depreciation of workers joining the booming technology sector. Acknowledging the existence of vintages within tasks-specific skills is important for two reasons. First, while the literature on tasks-specific skills stresses that task-specific skills are portable across firms and occupations (e.g., Gathmann and Schonberg, 2010), we emphasize that this portability might exist only within a given vintage of skills and erodes as new vintages appear. Second, skill vintages imply that different cohorts of workers can have different wage patterns, as observed for the ICT boom cohort of skilled workers relative to the post-boom cohort.

The existence of vintages among task-specific skills, demonstrated by a spike in the hiring of high-skill workers followed by a depreciation of their human capital, echoes the literature on investment booms and technological change showing that physical capital depreciates faster during booms (Jaimovich and Rebelo, 2009). Similar to physical capital, we show that the hiring boom induced by technological change translates into fastened depreciation of human capital.

Finally, our contribution differs from the classic result that the aggregate state of the economy has persistent effects on labor market entrants (Oyer, 2006; Kahn, 2010; Oreopoulos, Wachter, and Heisz, 2012; Altonji, Kahn, and Speer, 2016; Speer, 2016; Schoar and Zuo, 2017; Shu, 2016; Schwandt and Wachter, 2019; Nagler, Piopiunik, and West, 2020). We show that booms in the technology sector have long-term effects on labor market entrants joining the booming technology sector relative to same-cohort individuals joining other sectors, controlling for selection.

2 Sectoral Reallocation during the ICT Boom

2.1 Data

We use administrative data on French workers and firms. We describe here the main data sets used in the paper, and relegate the full list in Appendix A.

Linked employer-employee data are collected by the national statistical office based on a mandatory employer report of the gross earnings of each employee subject to payroll taxes. The data include all employed individuals in the private sector, with information on the gross and net wage, dated employment periods, number of hours worked, job occupation, and the individual's birth year and sex. The data also include unique firm and establishment identifiers that can be linked with other administrative data. The exhaustive employer-employee data do not include unique individual identifiers.

For a 1/24th subsample of the exhaustive employer-employee data (individuals born in October of even-numbered years), individuals are assigned a unique identifier that enables us to reconstruct their entire employment history. Individuals drop out of this panel data set only if they earn no wage in the private sector, because they exit the labor force, become unemployed, switch to self-employment and pay themselves only dividends, or move abroad.

We focus on the employer-employee panel over the years 1994–2015. Each observation corresponds to a unique firm-worker-year combination. We focus on job spells that are full time and last for at least six months in a given year. After we apply this filter, each individual has at most one job per year.[3] We obtain a panel at the worker-year level. Workers can have gap years in this panel when they earn no wage in the private sector, work part time, or over periods of less than six months.

The employer-employee data include a two-digit classification of job occupations that maps into the skill content of the job. We identify skilled workers as those holding higher-level occupations, which are comprised of "managers and professionals" (one-digit code 3) and "heads of company with at least ten employees" (two-digit code 23). Skilled workers represent 16% of the labor force over 1994–2015. Within managers and professionals, the two-digit classification distinguishes between occupations with a STEM skill content (two-digit code 38) and those with a management/business content (two-digit code 37), which represent 33% and 42%, respectively, of skilled jobs over 1994–2015, and heads of company with at least ten employees (code 23) represent another 4%.[4] Appendix Table B.1 reports summary statistics for the sample of skilled workers over the period 1994–2015. The median skilled worker is a man (fraction 69%), is 43 year old (mean 43), and earns an annual gross salary of 41,000 euros (mean 50,000 euros). Unless otherwise stated, all amounts in the paper are in constant 2000 euros. Finally, a 4/31 subsample of the employer-employee panel data (individuals born in the first four days of October) can be linked with census data, which contain demographics information. We use this smaller sample to retrieve information on education.

We retrieve information on firms from three sources. Firm accounting information is from tax files, which cover all firms subject to the regular or simplified corporate tax regime. Information on firm ownership structure is from a yearly survey of business groups run by the statistical office and crossed with information from Bureau Van Dijk. The data include information on both direct and indirect stakes and cross-ownership, which allows us to reconstruct group structures even in the presence of pyramids. The data include information on the nationality of the ultimate owner, which allows us to identify subsidiaries of foreign companies. Finally, we retrieve the list of all new business registrations with the event date from the firm register, and use this information to identify startups.

3. There are a few workers with full-time job spells of six months in two different firms in the same year. In these rare cases, we keep the observation with the higher wage.

4. The other two-digit occupations within managers and professionals are mostly for occupations held by self-employed or public sector workers: health professionals and legal professionals (code 31); public sector managers and professionals (33); teaching professionals (34); cultural professionals (35), which represent less than 1%, 8%, 9%, and 3%, respectively, of skilled jobs.

2.2 Labor Reallocation Over the ICT Boom-Bust Cycle

We analyze the late 1990s boom in the Information and Communications Technology (ICT) sector using the OECD (2002) definition of ICT industries. Appendix Table B.2 reports the list of four-digit ICT industries and their shares in total employment and in skilled employment during the sample period. The overall ICT sector represents 5% of total employment and 15% of skilled employment, reflecting that ICT is intensive in skilled labor. The fraction of workers holding a master's degree is 14% over all industries, whereas it is 30% in the ICT sector. The ICT sector is more specifically intensive in STEM skills: The fraction of skilled workers in STEM occupations is 35% across all sectors and 70% in the ICT sector.

Figure 1 illustrates the boom and bust cycle in the ICT sector in the late 1990s. While modest for total employment (Panel A), the ICT boom is evident for skilled workers (Panel B). The share of the ICT sector in total skilled employment features a sharp deviation from an increasing trend during the 1998–2001 period, with the share going from 12.5% in 1996 up to 16.5% in 2001 and down to 15% in 2005.

Panel C shows that the deviation from the trend is entirely driven by labor market entrants. The figure decomposes the ICT sector's share of skilled employment (plotted in Panel B) into the part made of workers who entered the labor force four years ago or less, and the part made of workers who have been in the labor force for five years or more. The latter exhibits an upward trend but shows no significant deviation. By contrast, the component representing young workers exhibits a sharp upward deviation from the trend during the ICT boom.

Panel C shows that the deviation from the trend is entirely driven by labor market entrants. The figure plots the ICT sector's share of skilled employment separately for workers who entered the labor force four years ago or less and for workers who have been in the labor force for five years or more. The latter exhibits a slightly upward trend but no significant deviation from trend. By contrast, among young workers, the ICT share exhibits a sharp upward deviation from the trend during the ICT boom.

Since sectoral reallocation induced by the boom mostly happens at labor market entry, we focus on skilled labor market entrants in the rest of the paper. We define the entry year in the labor market as the year in which individuals take their first full-time job, subject to the condition that they are no more than 30 year old at that time.[5] Appendix Table B.1 reports summary statistics for skilled

5. We drop individuals who are older than 30 at entry. The results are robust to using a cutoff at 35 year old. Since the panel data start in 1976, there is no risk of mismeasuring entry because it would have happened before the first year of data.

individuals entering the labor market over 1994–2005. The median skilled entrant takes her first job at the age of 26 (mean 26) and has an annual gross salary of 38,000 euros (mean 45,000 euros).

Panel D shows that the share of skilled labor market entrants starting in the ICT sector exhibits a sharp deviation from the trend during the 1998–2001 period. The ICT sector share of skilled entrants almost doubles from 17.5% in 1996 to 31% in 1999, before dropping down to 19% in 2004.

To summarize, we have two main facts regarding labor reallocation during the boom. First, the ICT boom induces a large sectoral reallocation of skilled labor, which happens almost exclusively through the sectoral choice of labor market entrants. During the boom, the ICT sector absorbs one-third of skilled labor market entrants. Therefore, the boom may have significant aggregate long-term effects depending on how it impacts the human capital accumulation of this cohort of workers. Second, the boom is sharply delimited over time, from 1997/8 to 2001, which allows us to define precisely the "ICT boom cohort" of workers, who enter the labor market during the ICT boom, together with the "pre-boom cohort" and the "post-boom cohort" of workers, who enter the labor market in the period right before and right after the boom, respectively.

In the rest of the paper, we study the effect of the initial sectoral choice of the boom cohort of workers on their human capital accumulation. In the next section, we develop a simple model which shows how this effect can be inferred from the long-run wage dynamics of the different cohorts.

3 Model

3.1 Setup

Time is discrete and horizon is infinite. At the beginning of each period, a mass one cohort of workers enter the labor market and choose in which sector $k = 1, 2$ to work. With a slight abuse of notation, let $E_{k,t}$ denote both the mass and the set of labor market entrants going to sector k in period t. In line with the evidence presented in Section 2.2 that sectoral reallocation occurs mostly through the sectoral choice of labor market entrants, we assume workers cannot switch sector after the initial sectoral choice made at the time of entry.[6] Worker i from cohort c in sector k supplies $H_{i,c,k,t}$ efficiency units of labor in period t. At the

6. The assumption of no sectoral mobility can be derived as a result if human capital accumulated on-the-job is sector specific (Rogerson, 2005).

end of each period, a fraction δ of workers of every cohort exit the labor market.

Human capital $H_{i,c,k,t} = \log(h_{i,c,k,t})$ has two components:[7]

$$h_{i,c,k,t} = \theta_{i,k} + h_{c,k,t}. \tag{1}$$

$\theta_{i,k}$ is a worker fixed effect reflecting ability for sector k. The distribution of $(\theta_{i,1}, \theta_{i,2})$ across workers is the same in every cohort, with mean zero. $h_{c,k,t}$ is a process driving post-entry human capital accumulation or depreciation given by:

$$h_{i,c,k,c} = 0, \tag{2}$$

$$h_{i,c,k,t} = h_{i,c,k,t-1} + dh_{c,k,t}, \qquad t > c, \tag{3}$$

where $dh_{c,k,t}$ is a shock to the period t-stock of human capital of individuals who work in sector k during period $t-1$. Human capital shocks follow the autoregressive process:

$$dh_{c,k,t} = dh + \rho_h(dh_{c,k,t-1} - dh) + \varepsilon^h_{k,t}, \qquad t > c, \tag{4}$$

where $\rho_h \in [0, 1)$, $dh_{c,k,c} = dh$, and $\varepsilon^h_{k,t}$ has zero mean. $dh_{c,k,t}$ has unconditional mean dh.[8] $\varepsilon^h_{k,t}$ is a human capital shock affecting all cohorts of workers in sector k in period $t-1$. It reflects on-the-job learning. It can also reflect changes in firm-specific human capital upon (unmodelled) job termination and within-sector job mobility. When $\rho_h > 0$, human capital shocks are serially correlated, implying that their effect builds up progressively over time.

Each sector $k = 1, 2$ employs labor to produce an intermediate good with constant returns to scale:

$$X_{k,t} = Z_{k,t} \sum_{c=-\infty}^{t} (1-\delta)^{t-c} \int_{i \in E_{k,c}} H_{i,c,k,t} di. \tag{5}$$

$Z_{k,t}$ is sectoral productivity and follows the autoregressive process $z_{k,t} = \rho_z z_{k,t} + \varepsilon^z_{k,t}$, where $\rho_z \in [0, 1]$ and $\varepsilon^z_{k,t}$ is a productivity shock with mean zero. The infinite sum in (5) is the efficient quantity of labor supplied in sector k in period t by all cohorts of workers $c = -\infty, \ldots, t$. The efficient quantity of labor supplied by cohort c is equal to the fraction of workers from cohort c who are still active,

7. Throughout the paper, we use lowercase letters to denote logs of uppercase variables.

8. dh can be different from zero to allow human capital to drift over the lifetime of workers. We assume $dh < -\log(1-\delta)$ to ensure that the aggregate supply of efficient labor in equation (5) remains bounded almost surely.

$(1-\delta)^{t-c}$, times the efficient quantity of labor supplied by workers from cohort c who started in sector k, $i \in E_{k,c}$.

The model allows for sectoral shocks that affect all cohorts similarly and for sectoral shocks that affect different cohorts differently. An example of the former is a positive sectoral productivity shock, $\varepsilon^z_{k,t} > 0$, which raises the productivity of all workers in sector k. An example of the later is a negative human capital shock to workers already in sector k, $\varepsilon^h_{k,t} < 0$, which weighs on the human capital of current workers but does not affect the human capital of future cohorts of workers. In practice, this can happen when technology changes and new cohorts of workers enter with up-to-date knowledge whereas old workers remain with older vintages of knowledge.

The final good is produced using the intermediate goods with CES:

$$Y_t = \Big(\sum_{k=1,2} A_k X_{k,t}^{\frac{\sigma-1}{\sigma}} \Big)^{\frac{\sigma}{\sigma-1}}, \tag{6}$$

where $\sigma > 1$, and $A_1^\sigma + A_2^\sigma$ is normalized to 1. The wage rate per efficiency unit of labor is determined by the marginal productivity of labor:

$$w_{k,t} = a_k + z_{k,t} - \frac{1}{\sigma}(x_{k,t} - y_t). \tag{7}$$

The wage of worker i is equal to her human capital times the wage rate in her sector in the current period. In log terms:

$$w_{i,c,k,t} = h_{i,c,k,t} + w_{k,t}. \tag{8}$$

Workers derive log utility over per-period consumption with discount factor $\beta < 1$, and consumption is equal to the current wage.

Workers have idiosyncratic preferences over their career choice. Worker i incurs a non-pecuniary cost γ_i if they choose sector $k = 1$. The distribution of γ_i across workers is the same in every cohort. Worker i from cohort c going to sector k obtains expected utility[9]

$$U_{i,c,k} = \sum_{t=c}^{\infty} \beta^{t-c}\, \mathbb{E}_c[w_{i,c,k,t}] - 1_{\{k=1\}}\gamma_i, \tag{9}$$

where $\mathbb{E}_c[.]$ denotes expectation conditional on beginning-of-period c information. Worker i chooses sector $k = 1$ if and only if $U_{i,c,1} > U_{i,c,2}$. Since expected human

9. The effect of workers' exit rate δ on expected utility is impounded in the discount factor β.

capital accumulation is the same in both sectors, the expected wage differential between the two sectors for a given worker depends on the expected wage rate differential and initial human capital differential: $\mathbb{E}_c[w_{i,c,1,t} - w_{i,c,2,t}] = \mathbb{E}_c[w_{1,t} - w_{2,t}] + \theta_{i,1} - \theta_{i,2}$. Therefore, the set of entrants in sector $k = 1$ in period c is:

$$E_{1,c} = \left\{ i : \gamma_i < \sum_{t=c}^{\infty} \beta^{t-c} \left(\mathbb{E}_c[w_{1,t} - w_{2,t}] + \theta_{i,1} - \theta_{i,2} \right) \right\}. \tag{10}$$

We assume that, when expected wages are equalized across sectors, the sectoral allocation of new workers is proportional to the sector weights in the production function, that is, the mass of $\{i : \gamma_i < (\theta_{i,1} - \theta_{i,2})/(1-\beta)\}$ is equal to A_1^σ.

Workers' sectoral choices depend on expectations of future wages. These choices and the resulting equilibrium outcomes do not depend on whether workers hold rational expectations or not. The only difference between both cases is that, if expectations are not rational, workers are systematically surprised by the realization of wages. Assessing whether workers' expectations are rational is outside the scope of this paper.

3.2 Equilibrium

We solve for a stationary equilibrium using a first-order approximation when productivity shocks and human capital shocks are small. Proposition 1 states that the equilibrium can be characterized in difference between sector $k = 1$ and sector $k = 2$, which we denote using the operator Δ, e.g., $\Delta w_t = w_{1,t} - w_{2,t}$. The state of the economy can be summarized by three variables: the (exogenous) sectoral difference in productivity, Δz_t, the (exogenous) sectoral difference in average human capital shock, $\Delta \overline{dh}_t$, and the (endogenous) sectoral difference in the efficient quantity of labor supplied by old workers, $\Delta \ell_t = \log(L_{1,t}) - \log(L_{2,t})$, where $L_{k,t} = \sum_{c=-\infty}^{t-1} (1-\delta)^{t-c} \int_{i \in E_{k,c}} H_{i,c,k,t} di$. We denote steady state values with *. The proof is in Appendix D.

Proposition 1 *At the stationary equilibrium:*

$$\Delta w_t \;\simeq\; \Delta w^* + w_z \,.\, \Delta z_t + w_\ell \,.\, \left(\Delta \ell_t - \Delta \ell^*\right) + w_h \,.\, \Delta \overline{dh}_t, \tag{11}$$

$$\Delta E_t \;\simeq\; \Delta E^* + E_z \,.\, \Delta z_t + E_\ell \,.\, \left(\Delta \ell_t - \Delta \ell^*\right) + E_h \,.\, \Delta \overline{dh}_t, \tag{12}$$

where $w_z \in (0,1)$, $w_\ell < 0$, $w_h \geq 0$, $E_z > 0$, $E_\ell < 0$, $E_h \leq 0$, *and* $\Delta \ell_t$ *evolves* *according to:*

$$\Delta\ell_{t+1} - \Delta\ell^* \simeq (1-\delta)dH \, . \left(\Delta\ell_t - \Delta\ell^*\right) + \ell_E \, . \left(\Delta E_t - \Delta E^*\right) + \Delta\overline{dh}_{t+1}, \quad (13)$$

where $\ell_E > 0$*, and* $\Delta\overline{dh}_{t+1}$ *is a weighted average of human capital shocks* $\Delta dh_{c,t+1}$ *across all cohorts* $c \leq t$*.*

Consider first the effect of a positive productivity shock in sector 1 relative to sector 2: $\Delta z_t > 0$. Higher productivity increases the demand for labor in sector 1. Since old workers cannot switch sector, sectoral reallocation takes place through the sectoral choice of labor market entrants. The wage rate increases in sector 1 relative to sector 2 ($w_z > 0$ in (11)) in order to induce more entry in sector 1 ($E_z > 0$ in (12)). Therefore, a positive productivity shock in the ICT sector in the late 1990s can explain the high entry rate (see Panel D of Figure 1) and the concomitant high wages (Figure 2) in ICT during the period.

Next, consider the effect of there being an excess mass of old workers in sector 1 relative to sector 2: $\Delta\ell_t - \Delta\ell^* > 0$. Higher labor supply lowers the wage rate in sector 1 ($w_\ell < 0$ in (11)), which reduces entry in sector 1 ($E_\ell < 0$ in (12)).

Finally, consider the effect of a positive human capital shock to old workers in sector 1 relative to sector 2: $\Delta\overline{dh}_t > 0$. If human capital shocks are persistent ($\rho_h > 0$), old workers are expected to become more productive in the future, increasing labor supply and reducing the wage rate in the future. This makes entry less attractive in the current period ($E_h < 0$), which pushes the current wage rate up ($w_h > 0$).

Equation (13) describes how the efficient quantity of labor supplied by old workers evolves over time. The first term on the RHS reflects that a fraction δ of old workers exit the labor market in each period, while those who do not exit experience an expected increase in human capital dH. Thus, the efficient quantity of labor by old workers mean reverts at rate $(1-\delta)dH$. The second term shows that entry of new workers adds to the stock of old workers ($\ell_E > 0$). The third term is a shock to old workers' human capital, which affects the efficient quantity of labor they supply. This shock is a weighted average of the shocks received by all cohorts of old workers.

4 Human Capital Depreciation

We now estimate the long-term value of human capital accumulated during the ICT boom by skilled workers who start in the booming ICT sector. Our identification strategy derives from the model presented in the previous section. It relies on comparing wages across sectors and cohorts to control for demand and selection effects. We present the empirical strategy and graphical results in Section 4.1. The empirical strategy allows us to incorporate additional controls and fixed effects to further tighten the identification. These refinements are presented in Section 4.2.

4.1 Main Results

We start from the wage equation. Combining equations (1) and (8), worker i from cohort c starting in sector k earns in period t the log wage

$$w_{i,c,k,t} = w_{k,t} + \theta_{i,k} + h_{c,k,t}. \quad (14)$$

The log wage depends on the wage rate in the sector $w_{k,t}$, the fixed type of the worker $\theta_{i,k}$, and human capital accumulated since entry $h_{c,k,t}$. We now show how to identify the human capital component by running wage regressions and progressively adding fixed effects for interactions between years, cohorts, and sectors.

4.1.1 Simple Difference: Across-Sector (Within-Cohort) Variation

For each cohort, we compare the wage of workers who start in the ICT sector relative to workers from the same cohort who start outside the ICT sector over time. To make this comparison, we estimate the regression at the individual-year level and run one regression for each cohort c:

$$\log(Wage_{i,t}) = \sum_{t'} \beta_t^c \, ICT_{i,0} \times (t = t') + \alpha_t \times X_{i,0} + e_{i,t} \quad (15)$$

$Wage_{i,t}$ is the annualized wage of individual i in year t. $ICT_{i,0}$ is a dummy equal to one if individual i starts in the ICT sector. It is interacted with cohort dummies and year dummies. The baseline specification includes cohort×year fixed effects $\alpha_{c,t}$, and their interactions with the vector $X_{i,0}$ of worker characteristics, which includes sex, age and age squared at entry, and entry year. $e_{i,t}$ is an error term clustered at the individual level.[10]

We estimate the regression separately for three cohorts: the *pre-boom* cohort, the *boom* cohort and the *post-boom* cohort, which are made of workers who start

10. We refer to this specification as a simple difference although it is formally a double-in-difference because of the time dimension (i.e., we compare the wage *dynamics* of two groups of individuals).

between 1994 and 1996, 1998 and 2001, and 2003 and 2005 respectively.[11] Figure 2 presents the estimated β_t^c for each cohort. Focusing first on the boom cohort, the figure shows that workers who start in the ICT sector during the boom earn an entry wage on average 5% higher than individuals from the same cohort and with the same characteristics, starting outside the ICT sector. However, the wage difference vanishes rapidly after the boom ends in 2001. Strikingly, the wage difference keeps falling after the bust. By 2015, workers who started in the booming ICT sector earn on average 6% less than workers from the same cohort who started outside the ICT sector.

The regression coefficient β_t^c captures the average wage difference in year t between workers from cohort c who start in the ICT sector and workers from the same cohort and with the same characteristics, who start outside the ICT sector. Equation (14) from the model shows that this difference encompasses three economic forces:

$$\beta_t^c = \Delta w_t + \Delta\overline{\theta}_c + \Delta h_{c,t} \tag{16}$$

where Δ denotes the difference operator between the ICT sector and other sectors and $c \in \{pre,\ boom,\ post\}$ is the cohort. Δw_t is the wage rate in year t in the ICT sector minus that outside the ICT sector. $\Delta\overline{\theta}_c$ is the average type of workers from cohort c who start in the ICT sector minus that of workers who start outside the ICT sector. $\Delta h_{c,t}$ is human capital accumulated from entry until year t by workers from cohort c who start in the ICT sector minus that of workers who start outside the ICT sector.

Equation (16) shows that β_t^{Boom} can fall over time for two reasons. First, there may be a secular decline in the wage rate in the ICT sector relative to other sectors (i.e., Δw_t decreases over time). For example, labor demand in ICT may persistently decline after the bust. Second, human capital accumulated by the boom cohort in the ICT sector may depreciate over time compared to human capital accumulated in other sectors (i.e., $\Delta h_{Boom,t}$ decreases over time). The selection term $\Delta\overline{\theta}_c$ however, cannot explain the drop in relative wages, as it is time-invariant and thus, can affect the level but not the dynamics of β_t^c. We further address selection in Section 4.1.3.

11. We include a gap year between each successive cohort to have sharply delimited cohorts. The results are robust to including the gap years in either one of the adjacent cohorts.

4.1.2 Difference-in-Differences: Across-Sector and Across-Cohort Variation

Human capital accumulated during the ICT boom ($\Delta h_{Boom,t}$) can be disentangled from labor supply and demand shocks (Δw_t) by comparing the wage dynamics of the boom cohort to that of the post-boom cohort. As shown by equation (16), labor supply and demand shocks Δw_t affect all cohorts equally. Therefore, we can absorb the labor supply and demand component by comparing the wage dynamics of the boom cohort relative to that of the post-boom cohort, and thus isolate the human capital component. Writing equation (16) in first difference between both cohorts:

$$\begin{aligned}\beta_t^{Boom} - \beta_t^{Post} &= \left(\Delta\overline{\theta}_{Boom} - \Delta\overline{\theta}_{Post}\right) + \left(\Delta h_{Boom,t} - \Delta h_{Post,t}\right) \\ &= \left(\Delta\overline{\theta}_{Boom} - \Delta\overline{\theta}_{Post}\right) + \sum_{\tau<2003} \frac{1-\rho^{t-\tau+1}}{1-\rho}\Delta\varepsilon_\tau^h, \end{aligned} \tag{17}$$

where the second equality follows from calculating $h_{c,t}$ using (2)–(4). The first term on the right side of (17) is time invariant and reflects selection, which we analyze below. The second term captures the long-run evolution of human capital accumulated in the ICT sector during the boom.

Figure 2 shows that, in sharp contrast with the boom cohort, the post-boom cohort shows no downward trend in the wage dynamics of workers who start in ICT. Therefore, the long-run wage discount of boom-cohort workers joining the ICT sector is not explained by a secular decline in labor demand or over-supply of skilled workers in the wake of the ICT bust. It is instead consistent with human capital accumulated during the boom depreciating over time.

The cross-cohort comparison can be made more rigorously by including a starting sector×year fixed effect (i.e., the $ICT_{i,0}$ dummy interacted with year dummies) in equation (15) and estimating the following equation:

$$\log(Wage_{i,t}) = \sum_c \sum_{t'} \beta_t^c\, ICT_{i,0} \times (i \in c) \times (t = t') + \alpha_{c,t} \times X_{i,0} + e_{i,t} \tag{18}$$

Including a starting sector×year fixed effect $\alpha_{c,t}$ implies that β_t^c is identified relative to a reference cohort. We use the post-boom cohort as the reference, i.e., as the control group. This specification is a difference-in-differences comparing the wage dynamics of workers who start in ICT relative to workers who start in other sectors (first difference) for workers of the boom cohort relative to workers of the post-boom cohort (second difference). The regression coefficients β_t^c are plotted in Panel A of Figure 3. The downward trend indicates a progressive depreciation of human capital accumulated by the boom cohort in the ICT sector during the boom.

4.1.3 Selection

We now study if our results are explained by negative selection into ICT during the boom, i.e., by the possibility that the booming ICT sector attracts a disproportionate share of low productivity workers.[12] As discussed above, the selection term in equation (17) is constant through time and thus cannot explain the downward trend in the boom cohort's wage dynamics. Moreover, the fact that the boom cohort and the post-boom cohort have similar average wages in 2003 (that is, right after the end of the boom and before the value of human capital starts to drift over time) suggests that the pool of workers in the boom cohort does not start from a particularly low level of human capital. That is, the selection term in (17) is likely small.

A starting sector×year fixed effect adequately controls for selection if unobserved heterogeneity shifts the wage profile by a time-invariant term, as $\theta_{i,k}$ does in the wage equation (14). However, a starting sector×year fixed effect fails to control for selection if unobserved heterogeneity is correlated with wage growth and not just with the wage level. In this case, the downward trend in the boom cohort wage dynamics could be explained by a subtler form of negative selection: the booming ICT sector might draw workers who would have experienced lower wage growth even if they had started in another sector.

We develop two strategies to control for selection correlated with wage growth. First, we augment the difference-in-differences specification with a large set of fixed effects interacted with time and cohorts to remove biases arising from selection of workers into different places, occupations and types of firms. We describe this strategy in detail in Section 4.2.

Second, we bring in the pre-boom cohort into the analysis. Individuals entering the labor market before the ICT boom experience the same human capital shocks and sectoral productivity shocks as individuals from the boom cohort. However, as shown in Panel D of Figure 1, the boom in the ICT sector was sudden, making it unlikely that anticipation of the boom led to negative selection among individuals

12. In the sectoral choice model laid out in Section 3, higher entry in a sector may increase or decrease the average type in the sector. Selection based on sector-specific type $\theta_{i,k}$ tends to decrease the average type. Selection based on non-pecuniary preferences γ_i can increase or decrease the average type depending on the joint distribution of non-pecuniary preference and types.

who started in ICT a few years before the boom. Therefore, if the long-term wage decline experienced by the ICT boom cohort is explained by negative selection during the boom, the pre-boom cohort should not experience it. By contrast, if the long-term wage decline is explained by depreciation of human capital accumulated in the ICT sector during the boom, the pre-boom cohort should also experience it.

As shown in Figure 2, the pre-boom cohort's wage dynamics has a downward trend very similar to that of the boom cohort. Panel B of Figure 3 estimated using starting sector×year fixed effects shows that the difference between the two cohorts is statistically insignificant.

4.2 Additional Fixed Effects

The difference-in-differences specification allows us to further deal with non-random allocation of workers into sectors, places and firms by conditioning on an additional extensive set of fixed effects.

The baseline difference-in-differences specification is equation (18) augmented with starting sector×year fixed effects. It compares the wage dynamics of workers going to ICT versus non-ICT, from the boom cohort relative to the post-boom cohort. The regression includes workers from the boom cohort and post-boom cohort. The β_t^c coefficients are identified for one cohort relative to the other cohort from year 2003 (the first year of the post-boom cohort) onward. We use the post-boom cohort as the reference group. For the sake of exposition, we do not include the triple-interaction terms $ICT_{i,0}\times$ cohort dummy $\times$ year dummy for every year from 2003 to 2015 but replace the year dummies with three time period dummies 2003–2005, 2006–2010 and 2011–2015. The specification is:

$$\log(Wage_{i,t}) = \sum_{\substack{\text{period=2003-05,}\\ \text{2006-10,2011-15,}}} \beta_{period}\, ICT_{i,0} \times BoomCohort_i \times (t \in period) + ICT_{i,0} \times \delta_t + \alpha_{c,t} \times X_{i,0} + e_{i,t} \quad (19)$$

The baseline specification controls for different worker demographic characteristics (age, age squared and sex) gathered in vector $X_{i,0}$, interacted with cohort×year fixed effects to account for potential changes in return to experience, changes in learning about the best occupational fit (Gervais, Jaimovich, Siu, and Yedid-Levi, 2014), or the gender gap. Workers may also select into occupations with different wage profiles. This could happen if the type of skills accumulated in the labor market (specific, general or portable) varies by occupation, for instance because the potential for on-the-job learning and human capital depreciation vary across occupations. We deal with this issue by including occupation×cohort×year fixed effects, which control for occupation-specific time-varying shocks.[13]

Workers might also sort across places based on their ex-ante productivity, exposing them to geographical differences in the potential for post-schooling human capital accumulation (e.g., Roca and Puga, 2016) and to various local labor shocks (e.g., Tuhler, Dustmann, and Schonberg, 2017; Gathmann, Helm, and Schönberg, 2018). We correct for this potential spurious correlation by adding commuting zone×cohort×year fixed effects.[14] This strategy controls for the fact that ICT workers or ICT firms might select into better places, and for any spurious correlation between ICT worker initial location and place characteristics.

Workers with different ex-ante productivity could be on different trends, or able to select across levels of occupations or places that are finer than our set of fixed effects. We address this potential problem by including entry wage quintile×year fixed effects, where the quintiles of entry wage are defined by year. In this case, we compare workers with the same initial wage.

With this set of fixed effects, the identification relies on comparing workers exposed to the same occupation and local labor market unobserved shocks, between ICT and non-ICT sectors and across cohorts. Workers or firms could still sort across specific ICT sub-sectors that differ in their future dynamics, potential for on-the-job learning (e.g., Parent, 2000) or level of unionisation and collective bargaining rules that could affect training programs and wage dynamics (e.g., Dustmann and Schönberg, 2009). We address this specific sorting by including four-digit sector×year fixed effects. Thus, we compare workers starting in the same narrowly defined sector from the boom cohort versus post-boom cohort. This strategy controls for selection of workers and firms across specific sub-sectors during the boom, and more generally for any change in the sub-sectoral composition of ICT firms' hiring during the boom.

Such strategy deals with sorting between workers and sectors based on sector time-invariant differences. This still leaves open the possibility that sectors are exposed to time-varying shocks that are not related to the long-run value of skills

13. We construct the fixed effect using the occupation in the first job rather than the current occupation because the current occupation is endogenous to human capital accumulation. For the same reason, all the other fixed effects described in this section and constructed using the commuting zone, four-digit sector, broad sector, and firm characteristics, are measured in the first job.

14. We define commuting zones as *départements*, which partition France into 99 areas. The results are also robust to using instead *bassins d'emploi*, which partition France into 380 areas.

acquired on-the-job, but still affect workers' wage dynamics differentially across cohorts. While we cannot add four-digit sector×cohort×year fixed effects because this is the level at which the variable of interest is defined, we can still partially deal with the issue by adding broad sector×cohort×year fixed effects. We use a classification with 32 broad sectors, four of which contain both ICT and non-ICT sub-sectors (Business Services; Manufacturing of Electric and Electronic Equipment; Transportation and Telecommunications; and Wholesale Trade). In doing so, we compare workers starting in an ICT four-digit sector to workers of the same cohort starting in the same broad sector, but in a non-ICT four-digit sector. For example, we compare "IT Consultancy" with "Management Consultancy"; or "Manufacturing of Computers" with "Manufacturing of Low-voltage Power Distribution and Control Systems"; but not "Manufacturing of Computers" with "Management Consultancy".

Despite the tight identification strategy, it could still be the case that within occupation-sector-local labor markets, high-quality workers sort into high quality firms. This could be a problem for identification since workers' long-term outcomes have been shown to be associated with firm characteristics such as size (e.g., Bloom, Guvenen, Smith, Song, and Wachter, 2018; Arellano-Bover, 2020), firm age (e.g., Ouimet and Zarutskie, 2014; Burton, Dahl, and Sorenson, 2017), and firm productivity (e.g., Abowd, Kramarz, and Margolis, 1999; Card, Heining, and Kline, 2013; Arellano-Bover and Saltiel, 2020).

Ideally, we would like to add fixed effects for the worker's initial employer interacted with year fixed effects, as in Cornelissen, Dustmann, and Schönberg (2017), and compare workers from different cohorts starting in the same firm. Because of the way the data are sampled, whereby individuals are randomly selected to be part of the panel irrespective of their employer, we cannot implement this strategy as few firms beyond the largest ones hire sampled high-skill workers from several cohorts. However, it is still possible to absorb a lot of the potential endogeneity coming from the correlation between firm characteristics and worker productivity by creating "pseudo firms" based on their characteristics and controlling for characteristics of the worker's initial employer interacted with year fixed effects. Given the importance of firm size, firm age and firm productivity emphasized in the literature, we use log employment, a startup dummy equal to one if the firm is two year old or less, and labor productivity measured as value added per worker.

Panel A of Table 1 reports the results when we progressively add the different fixed effects described above. Column 1 shows the baseline specification. The results are in line with Figure 3. During the 2003–2005 period, individuals from the boom cohort who started in ICT have similar wages as individuals from the post-boom cohort who also started in ICT (relative to the same comparison for individuals who started in other sectors.) However, as time goes by, individuals who started in the booming ICT sector experience slower wage growth such that their wage is 6.2% lower over the 2011–2015 period. We then add occupation×cohort×year fixed effects (column 2), local labor market×cohort×year fixed effects (column 3), entry wage quintile×cohort×year fixed effects (column 4), four-digit sector×year fixed effects (column 5), broad sector×cohort×year fixed effects (column 6), firm controls interacted with cohort×year fixed effects (column 7), and finally all of them together (column 8). Reassuringly, the estimates are stable across specifications. Starting in the ICT sector during the boom leads to between 6.2% and 11.0% average lower wages by 2011–2015. If anything, saturating the regression with fixed effects tends to increase the point estimate.

While holding fixed as many confounding factors as possible, the results might be contaminated by composition effects due to non-random attrition in the data. This would happen for instance if successful tech workers migrate to other countries such as the US or unsuccessful workers drop out of the labor force, and if the extent of attrition varies across cohorts. In Panel B, we control for composition effects by including individual fixed effects. Individual fixed effects ensure that we identify wage changes off individual wage trajectories and not off changes in the pool of workers induced by attrition. The inclusion of individual fixed effects imply that the β coefficients are identified relative to a reference period of time. We use 2003–2005 as the reference period. The coefficient for the period 2011–2015 in column 1 implies that workers starting in the booming ICT sector experience 6.0 percentage points slower wage growth from 2003–2005 to 2011–2015. This is in line with the results in Panel A (-0.062 for 2011–2015 relative to -0.008 for 2003–2005 corresponds to 5.4 percentage points slower wage growth).

While the specification with worker fixed effects guarantees that the estimates are not contaminated by a change in the composition of workers, differential attrition across cohorts could still bias the results if attrition is correlated with systematically better or worse wage trajectories, i.e., not just with the wage level but also with wage growth. In this case, the counterfactual wage that individuals would have earned if they had not dropped out of the data is on average different from that of individuals who do not drop out of the data even after controlling for worker fixed effects. This bias cannot be estimated directly but we can take a clue from the wage dynamics before individuals drop out of the data.

We define an exit dummy that equals one if the individual permanently exits

from the employer-employee data in the next year. The last year of data is 2015, so we define the exit dummy until 2010 to reduce truncation bias. We regress the exit dummy on the worker's wage growth over the past two years interacted with the ICT dummy and the boom cohort dummy, controlling for the same set of fixed effects as in equation (19). Results are reported in Table 2. In column 1, the negative coefficient on wage growth implies that workers who exit from the data tend to have slower wage growth on average. In column 2, the negative coefficient on wage growth interacted with $ICT_{i,0}$ implies that workers who started in ICT are on average more likely to exit the sample when they are on a growing wage trajectory. The key result is in column 3, showing that this relation is not specific to the boom cohort. The coefficient on wage growth interacted with $ICT_{i,0}$ and the boom cohort dummy is statistically insignificant and the point estimate is essentially zero. It implies that there is no differential pre-exit wage growth between workers who started in ICT during the boom relative to workers who started outside of ICT and relative to workers who started after the boom. Therefore, the results on the wage dynamics of the boom cohort of ICT workers are unlikely to be biased by variation in the determinants of attrition.

4.3 Robustness

Underperforming firms. While France fully embraced the ICT revolution and produced successful ICT firms, the country has not become the worldwide leader in that sector. As such, the wage discount might be specific to employees of French firms. To test this, we exploit the fact that many large US firms have offices across the world, including in France, so their employees located in France appear in our data. We use ownership data to identify subsidiaries of US companies defined as firms that are 100% owned by a US company.[15] In column 1 of Table 3, we re-estimate the baseline regression for this subset of firms and find a similar effect as on the entire sample of firms.

In a similar vein, one might suspect that the wage discount originates from ICT employers with little or mild commercial success. In column 2 of Table 3, we restrict the sample to workers taking their first job in a firm with sales growth over the next five years above 40% (the top quartile of the distribution). Here again, we find a similar negative point estimate as on the entire sample of firms. Therefore, it appears that the long-term wage discount is neither a French firm nor a low-quality firm phenomenon. Irrespective of the quality and type of employers,

15. Examples of US employers in the ICT sector in the data include Microsoft and IBM.

workers with human capital highly exposed to the development of ICT technologies during the boom experience accelerated depreciation of their human capital.

Measurement of earnings. Workers' earnings may be under-estimated because the employer-employee data report labor income but not capital income. Capital income can be significant for entrepreneurs. It may also be relevant for employees granted shares or options in their employer's stock. To account for capital income, we link the employer-employee data with employers' financial statements from tax filings. Since we do not have information on stock grants or stock options, we calculate capital income under two different assumptions. First, we assume the CEO holds all cash flow rights and add the firm's net income to the CEO's earnings.[16] Second, assuming employees have ownership stakes when they join startup companies, we allocate during the first eight years of a firm's life one-third of its net income to the skilled employees who joined the firm within three years of firm creation.[17] In both cases, we calculate workers' total earnings as wage plus capital income and use log of total earnings as the dependent variable. Column 3 of Table 3 reports the results when firms' profits are assigned to the CEO and column 4 when startup firms' profits are shared among skilled workers. In both cases, accounting for capital income has little effect on the magnitude of the long-run wage discount.

Education. The fact that the pre-boom cohort experiences the same long-run wage discount as the boom cohort provides strong evidence that selection is unlikely to drive our results. We can test directly for one important dimension of selection by testing whether the level of education of workers starting in the ICT sector during the boom stands out. We focus on the subset of individuals in the matched employer-employee data that can be linked with education information from Census data. We estimate equation (19) using as the dependent variable a dummy equal to one if the worker holds a master's degree.[18] Because education is

16. We identify the CEO as one-digit occupation code 2. When the firm reports several CEOs, we split the net income equally among them. Results are similar when we use dividends instead of net income. We prefer net income because it includes capital gains coming from undistributed profits.

17. We assume that this one-third fraction of net income is shared between the early joiners of the startup in proportion to their wage. We use a profit share of one-third because it is unlikely capital providers would not claim at least two-thirds of the profits (e.g., Eisfeldt, Falato, and Xiaolan, 2019). Results are robust are robust to using different profit shares and different time horizons at which we assume ownership stakes are granted to employees.

18. Master's degrees correspond to at least five years of higher education and include degrees from French elite *Grandes Ecoles*, university masters, and doctorates.

constant over time, we use only one observation per worker. Appendix Table B.3 reports the results. Across the different specifications, we find no evidence that the pool of workers going in the ICT sector during the boom is of lower quality based on their education achievements.

Quantile regressions. A career start in the booming ICT sector is associated with low average long-term wage growth. A possible interpretation is that such a career start exposes workers to high idiosyncratic risk because of the uncertainty regarding which firms and technologies will prevail in the long run. In this case, the low average wage growth may conceal a small probability of success, positive skewness, and high wage growth in the right tail of the distribution.

Table 4 reports estimates of quantile regressions for the 10th, 25th, 50th, 75th and 90th percentiles of wage growth. The long-term wage discount experienced by individuals starting in booming ICT sector is fairly uniform across the wage growth distribution, ranging from 7.1% (at the 90th percentile) to 9.0% (at the 10th percentile). This result rules out the interpretation that the average discount is associated with a small probability of very positive outcomes. Thus, the boom does not create winners and losers among skilled individuals who joined the booming ICT sector, but instead shifts their entire wage growth distribution to the left.

Cumulative earnings. The long-term wage would not accurately reflect long-term productivity if there was reverse backloading, i.e., if workers earned high upfront wages in exchange for lower wages later on (Lazear, 1981). In this case, individuals starting in the booming ICT sector might still earn the same cumulative earnings as individuals starting in other sectors despite slower wage growth. To test whether this is the case, we re-estimate equation (19) using as the dependent variable cumulative earnings (including from part-time and short job spells, which were excluded from the previous regressions) from labor market entry up to each year t post-entry, discounted back to the entry year at a rate of 5% per year. We do not include individual fixed effects because we want to analyze cumulative earnings as a stock, not as a change relative to a reference period.

Studying cumulative earnings in log (column 1 of Table 5), we find that skilled workers starting in ICT during the boom earn cumulative earnings from entry to 2015 that are 4.1% (significant at 5%) lower than that of similar workers starting in other sectors. When we analyze cumulative earnings in level (column 2), the discounted cumulative earnings loss is 19,600 euros (significant at 1%). Column 3 shows that this estimate is robust to accounting for unemployment benefits.[19]

5 Explaining Human Capital Depreciation

The results in the previous section show that the relative wage of workers who started in ICT during the ICT boom declines over time, and that this relative decline is neither explained by selection nor by labor market imbalance. Based on the wage decomposition (14) implied by our model, we conclude that the relative wage decline is due to the depreciation of these workers' human capital. Two mechanisms can explain such accelerated depreciation of skills. First, skills associated with tasks embedding the rapidly changing information and communications technologies may become obsolete. This leads to the co-existence of several vintages of skills, with old vintages losing value over time as technology changes (Chari and Hopenhayn, 1991). Second, skills may have limited portability across firms due to the loss of firm-specific human capital (Becker, 1975), adverse selection (Gibbons and Katz, 1991), and search frictions. If the boom-bust cycle in the ICT sector leads to a higher probability of job termination for the boom cohort, limited skill portability might explain the wage discount.

5.1 Skill Obsolescence

If the wage discount is explained by the obsolescence of skills associated with the rapidly changing ICTs, wages should decline faster for workers performing tasks that are more directly associated with the production of these technologies. We test this hypothesis in two ways.

First, we exploit the fact that the occupation classification in the data is organized by broad levels of qualification, which allows us to classify workers into high-, middle- and low-skill workers. We proxy the technological content of human

19. Since unemployment benefits (UB) are only reported starting in 2008, we assign estimated UB when an individual has no earnings reported in the data in a given year. In France, individuals are entitled to UB if the job is terminated or not renewed by the employer, but not if they resign, and UB are paid for a period of time roughly equal to that of their pre-unemployment job spell and no longer than two years (Cahuc and Prost, 2015). Since the data do not report the motive for job termination, we assume in the baseline scenario that all job terminations give rise to one year of UB equal to the average replacement rate in France of 60% of the total wage earned in the previous year. We obtain an UB-adjusted cumulative earnings loss that varies within a range of 500 euros of that of the baseline scenario when we use a more conservative replacement rate of 30% to account for the fact that not all job terminations give rise to UB, and when we use a more aggressive UB length of two years if the pre-unemployment job spell lasts for at least two years.

capital with the level of the worker's qualification and assume that high-skill workers' human capital embeds more technology-specific skills and is thus more prone to obsolescence when new technologies appear than that of middle-skill workers and even more so than that of low-skill workers. This assumption is consistent with results in Gathmann and Schonberg (2010), who find evidence that high-skill workers accumulate more task-specific human capital on the job.

In Table 6, we re-run the baseline regression (19) separately for the three groups of workers: high-skill workers (i.e., the same sample as in the previous section), middle-skill workers, and low-skill workers. We find that high-skill workers starting in the ICT sector during the boom experience a larger long-run wage discount (column 1) than middle-skill workers (column 2), while the effect is insignificant for low-skill workers (column 3), consistent with a specific obsolescence of more advanced skills associated with the new and rapidly changing technologies.

This result could be partly explained by the fact that middle-skill and low-skill workers are exposed to different shocks than high-skill workers that are independent from human capital considerations. Our second test addresses this issue by exploiting variation *within* high-skill workers in the degree to which they have been exposed to new technologies because of the firm or sector they work in. We construct two proxies of the job's technological content of high-skill workers. The first one measures the technological intensity of the individual's first employer, defined as the fraction of high-skill workers in the firm's workforce. The second proxy measures the technological intensity of specific (four-digit) ICT sectors in which the individual starts her career. It is defined as the fraction of high-skill workers in the four-digit sector.[20] We re-estimate the baseline specification (19) on high-skill workers and we interact the boom cohort and time period dummies with the variable $HighTechICT_{i,0}$ equal to the worker's job technological content, as measured by one of the two proxies, times $ICT_{i,0}$.

Results are in Table 7. In column 1, the coefficient on $HighTechICT_{i,0} \times BoomCohort \times$ 2011-15 is negative and significant at the 1% level. Thus, the wage discount is larger for workers who started in more technology intensive ICT firms. The point estimate implies that the wage discount is 4.9 percentage points larger for a worker starting at an ICT firm at the 25th percentile of the technological intensity distribution than at the 75th percentile (inter-quartile range is 0.28). This finding might reflect a more general pattern by which skilled workers start-

20. The top three ICT sectors in terms of technological intensity are "IT consultancy", "Software", and "Other IT-related activities", while the bottom three are "Manufacturing of insulated wires and cables", "Manufacturing of capacitors", and "Manufacturing of office devices except computers".

ing in more technology intensive firms even outside ICT would experience slower long-term wage growth. To control for this, in column 2, we include the firm's technological intensity outside the ICT sector (i.e., non-interacted with $ICT_{i,0}$) interacted with the boom cohort and time period dummies. We find that the impact of the firm's technological intensity for workers starting in ICT remains significant (albeit at the 10% level) and the point estimate remains large.

A similar pattern emerges when we use the proxy for the sector's technological intensity. Column 3 shows that the long-term wage discount is larger for workers who started in more technology intensive sectors within the broad ICT sector. The point estimate implies that the wage discount is 6.5 percentage points larger for a worker starting in an ICT sector at the 25th percentile of the technological intensity distribution than at the 75th percentile (inter-quartile range 0.20). Column 4 shows that this finding is not explained by workers starting in more technology intensive sectors even outside ICT experiencing slower wage growth. The sector's technological intensity has no significant effect on the long-term wage growth of workers starting during the boom outside ICT (the coefficient on $HighTech_{i,0} \times BoomCohort \times$ 2011-15 is small and statistically insignificant), which is consistent with rapid obsolescence of technical skills acquired specifically in the ICT sector during the boom, but not with a general trend of obsolescence of technical skills in the rest of the economy.

5.2 Job Losses

Part of the decline in relative wages might be explained by the fact that the skills workers acquire on-the-job have limited portability, for instance because these skills are firm specific (Becker, 1975). This would be an issue for high-skill workers who start in the ICT sector during the boom if they experience a higher probability of job termination in the bust. For this explanation to explain our result, two facts must hold true: (1) starting in the ICT sector during the boom is associated with a higher rate of job termination, and (2) this higher rate of job termination explains the wage discount in the wage regression.

We start by testing if workers who start in the ICT sector during the boom are more exposed to displacement risk. We construct three measures of job termination: (1) a dummy variable equal to one if the worker experiences a job termination within the first four years after entry; (2) a dummy equal to one if the worker experiences a job termination within the first four years and employment at the worker's employer declines by more than 10% in the year of the job termination; and (3) a dummy equal to one if the worker experiences a job termination within the first four years and the transition to the next job leads to a wage cut for the worker.

While the first measure uses all instances of job termination, the last two ones are designed to capture forced job termination, which is exogenous to the worker's human capital.[21] We construct these three measures of job termination regardless of whether the next job is in the same sector or not (columns 1 to 3 of Table 8) as well as by conditioning on the fact that it also leads to a change of four-digit sector in the next job (columns 4 to 6). We collapse the data at the worker level because the measures of early career job termination are time-invariant. We regress job termination on the same set of explanatory variables as in our baseline specification (but without the time dimension).

In column 1 of Table 8, we find that while workers who start in the ICT sector are more likely to change job in general, there is no differential effect for the cohort that starts during the boom relative to the cohort that starts during the post-boom period. The coefficient on $ICT_{i,0} \times Boom_i$ is not only statistically insignificant but the point estimate is essentially zero. When we focus on forced job termination in columns 2 and 3, a clearer pattern emerges. Workers starting in ICT during the boom are more likely to experience forced termination than workers starting in ICT after the boom. In columns 4 to 6, we replicate the analysis for mobility across sectors, as being forced to move across sectors may be associated with a loss of not only firm-specific human capital but also sector-specific human capital. We find that while the point estimates for workers who started in ICT during the boom are positive, none are statistically significant at conventional levels.

While high-skill workers starting in the booming ICT sector do face higher risk of job termination, a back-of-the envelop calculation shows that this is unlikely to explain the wage discount. Explaining the 7% wage discount (Panel B of Table 1) with the 4 percentage point increase in the probability of job termination (columns 2 and 3 of Table 8) would require that job termination leads to a $0.07/0.04 = 175\%$ decline in the log wage. This is several orders of magnitude larger than estimates in the literature. For instance, Gibbons and Katz (1991) find a 6.4% decline and Dustmann and Meghir (2005) find a 6% decline for skilled workers.

Nonetheless, in Table 9, we test formally whether the higher probability of

21. The second measure, which conditions on negative employment growth at the initial employer, is similar in spirit to using plant closure to generate exogenous events of job displacement (e.g., Gibbons and Katz, 1991, Dustmann and Meghir, 2005).

job termination explains the long-term wage discount by re-estimating the wage regression and directly controlling for job termination. We use our second measure of job termination, when the initial employer experiences at least a 10% reduction in its workforce, as this is the situation where job change is most likely exogenous and not mechanically related to future wages. Compared to the baseline long-term wage discount of 7.7% (column 1), controlling for job termination barely changes the point estimate, which decreases to 7.6% (column 2). Job termination during a sectoral bust might have a disproportionate impact on long-term earnings. To account for this possibility, in column 3, we include the interaction term between job termination and $ICT_{i,0} \times BoomCohort$ to allow job termination to have a different effect on workers starting in the booming ICT sector than on workers starting in other sectors. The coefficient on $ICT_{i,0} \times BoomCohort \times$ 11-15 is the long-term wage difference between workers starting in the ICT sector and experiencing no job termination, and entrants in other sectors experiencing no job termination. This long-term wage difference remains similar (7.2%) to the baseline (7.7% in column 1). Using job changes across sectors instead of across employers yields the same conclusion (columns 4 and 5).

6 Concluding Remarks

A popular argument holds that technology bubbles can be growth-enhancing because they promote investments that increase future productivity. This argument is formalized in speculative growth models such as Olivier (2000) and Caballero, Farhi, and Hammour (2006). We test a specific mechanism by which a bubbly technology sector can affect future productivity: accumulation of human capital by the large cohort of workers hired in the bubbly technology sector. We find no evidence for this mechanism, but instead find evidence for the opposite mechanism. The bubbly technology sector grows by hiring young skilled workers and paying them a wage premium as long as the bubble lasts. Fifteen years out, these workers have significantly lower wages than both same-cohort workers in other sectors and next-cohort workers in the technology sector. The long-term wage discount is not explained by negative selection or job losses in the bust. It hits harder workers holding higher-skill jobs and workers in skill-intensive firms and sectors, consistent with obsolescence of skills acquired during the bubble. To be clear, we do not suggest that bubbles cause skill obsolescence. Instead, technology bubbles are both the consequence of accelerating technological change (Shiller, 2000), which is the actual driver of skill obsolescence, and the cause of the large flow of young

skilled workers to the technology sector. Technology bubbles can therefore have a long-term effect on productivity by distorting the sectoral allocation of labor in a way that is adversely correlated with skill obsolescence.