

The Long-Term Consequences of the Tech Bubble on Skilled Workers' Earnings

Johan Hombert*

Adrien Matray[†]

June 4, 2018

Abstract

We use French matched employer-employee data to track skilled individuals entering the labor market during the late 1990s tech bubble. The boom led to a sharp increase in the share of skilled entrants in the tech sector, which offers relative higher wages at the time. When the boom ends, however, the wage premium reverses and these skilled workers end up with a 5.5% wage discount ten years out, relative to similar peers who started in a non-tech sector. Other moments of the wage distribution of the boom, pre-boom, and post-boom cohorts are inconsistent with explanations based on a selection effect or a cycle effect. Instead, we provide suggestive evidence that workers allocated to the booming tech sector accumulate human capital early in their career that rapidly becomes obsolete.

*HEC Paris and CEPR, hombert@hec.fr

[†]Princeton University, amatray@princeton.edu

1 Introduction

Radical technological change often comes with episodes of boom and bust in the technology sector, during which high stock prices and high investment in labor and capital are followed by sharp reversals (Schumpeter (1942)). During the boom, the tech sector attracts talents, and in particular young talents as sectoral choices are mostly made early in one’s career and young firms and sectors are inherently more likely to hire young workers (Ouimet and Zarutskie (2014)). These early career choices are determinant for skilled workers’ human capital accumulation and long-run productivity.¹ Thus, the large flows of young talents in and out of the tech sector during a boom-bust episode may have large and long-lasting consequences for productivity through its effect on skilled workers’ human capital. These consequences are beyond the direct effect of developing new technology: they determine the difference for long-run productivity between a smooth technological change process and the more bubbly process that characterizes technological change in the data.

How the allocation of talents to the booming tech sector affects their post-schooling accumulation of human capital and long-run productivity remains however an open empirical question. Starting in the booming tech sector may benefit skilled workers by exposing them to new technologies and enabling them to develop skills that may be redeployed in other firms or other sectors. This will enhance their productivity in the long run, even after the tech sector contracts when the boom ends. In this case, technological bubbles can have positive and radically different effects on workers’ productivity than low-tech sectors bubbles, such as housing bubbles as studied by Charles, Hurst, and Notowidigdo (2015).

On the other hand, skilled workers going into the booming tech sector may face a risk of human capital depreciation in the long run for at least two reasons. First, as the boom ends and the sector contracts, many firms downsize or go bankrupt, causing workers employed by these firms to lose job-specific human capital. Tech firms financed during booms may also have a higher probability of failure (Nanda and Rhodes-Kropf (2013)), exposing workers to a higher risk of investing human capital into a losing firm or technology. Second, technology can change rapidly, exposing talents attracted by the booming tech sector to obsolescent skills. Moreover, because the boom relaxes credit constraints (Brown, Fazzari, and Petersen (2009)), technologies developed during this period may be of lower average quality due to reduced project screening, or less radical to due changing firms’ incentives over the business cycle (Manso, Balsmeier, and Fleming (2017)).

Studying the long-run effects of boom-bust episodes in tech sectors on skilled workers’ productivity requires to be able to track individuals over a long period of time with information about jobs and wages. To do so, we use administrative matched employer-employee data that can be linked to firm balance sheet information for France over the period 1994 to 2010. Armed with these data, we

¹See Gibbons and Waldman (2006) for a model and Altonji, Kahn, and Speer (2016) for empirical evidence and additional references.

analyze the late 1990s boom-bust cycle in the Information and Communication Technology (ICT) sector and shed light on the following questions. First, what is the long-run wage dynamics of talents allocated to the tech sector during the boom? Second, can we attribute any difference in long-run productivity between these workers and otherwise similar individuals allocated to other sectors to a treatment effect of the technology boom—or do selection effects or other confounding factors prevent any inference? Third, which economic mechanisms can explain the long-run wage dynamics of talents allocated to the booming tech sector?

We start by documenting the impact of the ICT sector boom for labor market allocation. The share of skilled workers in the ICT sector sharply deviates from its long-run trend from 1997/8 to 2001, in line with the timing of the boom in the U.S. for R&D (Brown, Fazzari, and Petersen (2009)) and stock prices (Ofek and Richardson (2003)). This deviation from the trend is almost entirely explained by skilled workers who recently entered the labor market, consistent with the broader evidence that young workers are unconditionally more mobile (Topel and Ward (1992)). The share of skilled labor market entrants starting in the ICT sector almost doubles between 1996 and 1999, from 18% to 33%, before dropping back to 20% in 2004. The sharp delimitation in time of the boom-bust episode enables us to study wage patterns for relatively well identified cohorts of skilled workers entering the labor market either before, during, or after the boom.

To answer the first question, we focus on the boom cohort (1998–2001) and compare the wage dynamics of skilled workers starting in the ICT sector during the boom to that of otherwise similar individuals of the same cohort but starting in a different sector. We estimate a series of panel regressions, in which future wages are explained by the allocation to the ICT sector at the time of the boom. The estimation is made possible by the panel nature of the data, that allows us to relate employment and wage outcome of a given worker in the future to her characteristics and choices made when she enters the labor market.

Our main finding is that by 2010, ten years after the peak of the boom, skilled workers who started in the booming ICT sector had experienced average wage growth 8.5 percentage points lower than that of individuals starting in other sectors, after controlling for age, sex, broad occupation, and education. Only one third of this poor long-term wage growth is explained by the fact that the starting wage during the boom was higher in the ICT sector. During the quick contraction, workers who started in the ICT sector experience a progressive decline in their relative wage such that, by 2004, their wage advantage has disappeared. Such a pattern would be consistent with any sectoral fluctuation, but remarkably, the relative wage of these workers keeps declining almost monotonously after 2004 such that, by 2010, they end up earning 5.5% less than workers who started outside of the ICT sector. This result is quantitatively robust to including worker fixed effects to account for time-invariant unobserved heterogeneity across workers, to excluding entrants in the financial sector which experienced high wage growth during the 2000s, to controlling for local economic trends, and to controlling for educational attainment for the subsample of workers we can match with census

data.

Quantile regressions further show that the entire wage growth distribution of entrants in the ICT sector during the boom shifts left, compared to entrants in other sectors. The evidence is thus inconsistent with the technology boom creating winners and losers among talents who go into the booming ICT sector. In particular, skilled workers' returns from starting in the booming tech sector do not exhibit the same skewed distribution as startups' returns (Kerr, Nanda, and Rhodes-Kropf (2014)).

Regarding the second question, we provide several pieces of evidence consistent with a treatment effect of the ICT boom on skilled workers. We start by showing two pieces of evidence inconsistent with a selection effect by which the marginal workers attracted by the boom would be of low intrinsic productivity and would have experienced poor wage growth if even she had started in another sector. First, such a selection effect would induce a worsening of the quality of workers at the low end of the distribution, generating a larger drop in the bottom quantiles of wage growth than in the top quantiles. The quantile regressions reject this prediction of the self-selection mechanism.

To further separate a treatment effect of the boom from a selection effect, we exploit the cohort of workers who started in the ICT sector just before the boom (1994–1996) and thus experienced the boom. To the extent that the boom was not anticipated in the mid-1990s, these slightly older workers were treated by the boom but not selected. Thus, the difference in outcomes between the boom cohort and the pre-boom cohort can be interpreted as the selection effect of the boom, whereas the commonalities in outcomes can be interpreted as the treatment effect on skilled workers exposed to the boom. Consistent with an absence of selection for the pre-boom cohort, we find that pre-boom entrants in the ICT sector have similar wages as entrants in other sectors until 1997. While they experience a relative wage increase during the boom, their wage dynamics after the boom follows the same pattern as the boom cohort, ending up 7% below that of entrants in other sectors. The similar long-term wage dynamics of the pre-boom cohort and the boom cohort is more consistent with a treatment effect on workers experiencing the boom in the ICT sector than with a selection effect by which low quality workers would select into the booming sector.

The second alternative explanation is that the long-run wage decline reflects a cycle effect, by which the swift boom-bust episode the ICT sector experienced at the end of 1990s would have been purely a bubble with no valuable technological content and characterized by a decline in productivity post boom. While at odds with the accumulated evidence about this sector, such a story might explain the relative wage pattern in the data. We rule out this possibility by focusing on the post-boom cohort (2003–2005). We find that workers of this cohort who start in the ICT sector face lower entry wages than entrants in other sectors, but they progressively catch up over time. Therefore, while the relative wages of pre-boom and boom cohorts of ICT sector entrants keep falling after the boom ends, the pattern is reversed for ICT sector entrants who did not experience the boom.

Finally, we turn to the third question and explore two possible mechanisms to explain the decline in relative wage growth. First, during the bust, ICT firms are more likely to downsize or go bankrupt. Workers losing their jobs may lose job-specific human capital in the process and end on a different career path associated with long-term earnings losses.² Second, skills accumulated in the ICT sector during the boom may become rapidly obsolete as a result of technological acceleration (Violante (2002)), reducing the relative long-run productivity of workers exposed to the boom.

We find little support for the job termination mechanism. First, when we decompose wage growth the labor market entry during the boom to 2010 into a within job spells component and a between job spell component, we find that almost all of the relative wage decline happens within jobs. Second, controlling in the wage growth equation for different measures of forced job termination, we find that it explains at best a very small fraction of the relative wage decline. Thus, unconditional job termination does not seem to explain the wage growth discount. A more refined version of the job termination mechanism is that losing one's job in a shrinking sector is particularly detrimental to long-run wage growth. We account for this possibility by allowing the effect of job termination to vary across sectors, and obtain a similar result in this case.

Finally, we explore if the poor long-run wage growth can be explained by the rapid obsolescence of skills developed early in their career by workers starting in the booming ICT sector. We hypothesize that skill obsolescence should be an increasing function of the intensity of the worker's job technological content. We construct three proxies of job technological intensity. First, we distinguish among skilled workers between scientists/engineers and those holding management/business positions. Consistent with technical skills being more subject to obsolescence, we find that the long-run wage discount of starting in the booming ICT sector is concentrated on engineers, whereas workers starting as managers in the ICT sector experience only a small and insignificant decline. This result is not explained by engineers experiencing slower wage growth independently of their starting sector. The second proxy is the fraction of scientists and engineers over total employees in the firm employing the individual at the start of her career. The third proxy is the fraction of scientists and engineers at the four-digit sector level. In both cases, the long-run relative wage discount is stronger for workers who started in firms and sectors that are more tech-intensive within the ICT sector. Overall, these results are consistent with a mechanism where exposure of young workers to the tech sector in a period of rapid technological change can be detrimental to their long-term human capital accumulation.

Our paper contributes to several strands of literature. First, we contribute to the literature studying the real effects of the late 1990s tech bubble. Most of the literature has focused on the potential upside for listed firms. Brown, Fazzari, and Petersen (2009) show that high stock valuations allow high-tech firms to issue equity and invest and R&D, while Campello and Graham

²See Jacobson, LaLonde, and Sullivan (1993) for early evidence on long-term earnings losses caused by job displacement and Jarosch (2015) for a more recent treatment.

(2013) find that similar effects are at work for non-tech firms. Dong, Hirshleifer, and Teoh (2017) show that high stock valuations encourage firms to spend more in R&D. Regarding the implications for skilled workers, Choi, Lou, and Mukherjee (2017) show, using episodes of high stock prices across various sectors, that they distort college major choices. We complement this literature by looking at the long-term effect on workers exposed to this technology boom and showing that while it may promote firm investment, the boom also has negative long-term implications for workers exposed to it as they end up accumulating human capital that rapidly becomes obsolete.

Second, our paper relates to the literature studying the effect of aggregate fluctuations on the long-term earnings of labor market entrants. This literature has found that workers graduating in a recession earn persistently less than those graduating nearby peaks (see Kahn (2010), Oreopoulos, von Wachter, and Heisz (2012)). Devereux (2002) argues that this pattern is explained by the fact that workers graduating in a recession are matched to lower-level starting jobs, which causes them to lag behind in term of task-specific human capital for more important jobs years later, producing persistent earnings differences across firm-entry cohorts (Beaudry and DiNardo (1991), Baker, Gibbs, and Holmstrom (1994)). We complement this literature by showing that booms in specific sectors can also lead to substantial long-term relative earnings differential within a cohort, because some entrants accumulate post-schooling skills that will rapidly be outdated.

Our paper is closer to two recent papers focusing on specific sectoral booms and busts. Charles, Hurst, and Notowidigdo (2015) study the 2000s housing bubble in the US and show that it led to an increase in high-school drop out. Beaudry, Green, and Sand (2016) show using cross-sectional data for the US that around the year 2000, the demand for cognitive tasks underwent a reversal, which might have amplified unemployment after the 2008 crisis.

The rest of the paper proceeds as follows. We present the data in Section 2 and stylized facts on the ICT boom in Section 3. We analyze wage dynamics in Section 4, present evidence in favor of a treatment effect of the boom in Section 5, and explore economic mechanisms in Section 6. Section 7 concludes.

2 Data

2.1 Matched Employer-Employee Data

Exhaustive employer-employee data. We use several administrative data sources from France. First, we use matched employer-employee data (*Déclaration Annuelle des Données Sociales*) 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 includes all employed persons in the private sector, with information about the (gross and net) wage, the first day and last day of the job spell, the number of hours worked, the job occupation in a two-digit classification, and the individual's birth year and sex. The data also includes information about the employer, in particular unique firm and

establishment identifiers that can be linked with other administrative data sources. We shall refer to this data set as the exhaustive employer-employee data, which we use to study the composition of the labor force at the aggregate level, sector level, and firm level. The exhaustive employer-employee data does not, however, include unique individual identifiers.

Employer-employee panel. 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. We refer to this subsample of the data as the panel employer-employee data, which we use to study the wage pattern of individual workers over time. An individual exits the panel data only if she earns no wage in the private sector, because she drops out of the labor force, becomes unemployed, switches to self-employment and pays herself only dividends, or moves abroad.

We focus on the panel employer-employee data over the years from 1994 to 2010. Each observation corresponds to a unique firm-worker-year combination. In most of the analysis, 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 the panel when they earn no wage in the private sector or work part time or over periods of less than six months. We shall refer to this data set as the *full-time panel*.

Because the ICT sector is essentially composed of skilled workers and given the crucial role these workers play for long-term growth, most of our analysis focuses on skilled workers, which we identify based on the occupation held. The data includes a two-digit classification of job occupations (*Professions et Catégories Socioprofessionnelles* that strongly maps 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). They represent 16% of the labor force over 1994–2010. Within managers and professional, the two-digit classification distinguishes between occupations with a science/engineering 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 1994–2010, and heads of company with at least ten employees (code 23) represent another 4%.⁴ Table 1, Panel A reports summary statistics for the sample of skilled workers over the period 1994–2010. The median skilled worker is a man (mean 0.72), is 42 years old (mean 43), and earns an annual net wage of 33,750 euros, which corresponds to a gross salary of about 58,000 euros (mean 41,150 euros or about 71,000 euros gross; unless otherwise stated all amounts in the paper are in constant

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

⁴The 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 (code 33); teaching professionals (34); cultural professionals (35), which represent less than 1%, 8%, 9% and 3%, respectively, of skilled jobs.

2000 euros).

Census-matched subsample. Finally, a 4/30th subsample of the panel data (individuals born in the first four days of October) can be matched with census data (*Echantillon Démographique Permanente*), which contains demographics information. We use this smaller sample to retrieve information about education and conduct robustness checks to investigate whether the results in the main sample are driven by unobserved heterogeneity in the level of education.

Labor market entrants. Most of our analysis will focus on workers who enters the labor force over the period 1994–2005. We define the year of entry as the first year in which the individual enters into the full-time panel, subject to the condition that she is no more than 30 years old at the date of entry year.⁵ We focus on skilled entrants defined as individuals holding a higher-level occupation in their first job. We construct the sample based on the skill-content of the occupation at the time of entry rather on the skill-content of the current occupation, because the latter is endogenous to human capital accumulated during the individual’s career.⁶ Table 1, Panel B reports summary statistics for the period 1994–2010 for skilled individuals entering the labor market over 1994–2005. The median skilled workers takes her first job at the age of 26 (mean 26) and has a annual net wage of 28,300 euros or 49,000 euros gross (mean 32,500 euros or 56,000 euros gross).

2.2 Firm Data

Tax files. We retrieve accounting information about firms from tax files (*FICUS*), which cover all firms subject to the regular corporate tax regime (*Bénéfice Réel Normal*) or the simplified corporate tax regime (*Régime Simplifié d’Imposition*). Firms with annual sales below 32,600 euros (81,500 euros in retail and wholesale trade) can opt out and choose a special micro-business tax regime (*Micro-Entreprise*), in which case they do not appear in the tax files. Since the micro-business tax regime does not allow firms to deduct expenses and in particular wages from taxable income, it is mainly used by firms with no employees.

Firm creation. We use business creation files, which contain the list of all business creation with the date of registration. We use this information to construct firm age.

3 The ICT boom and Bust

3.1 The Information and Communications Technology Sector

We analyze the boom in the Information and Communications Technology (ICT) sector that took place in Europe and the U.S. at the end of the 1990s. We use the definition of ICT industries

⁵Since the panel data starts in 1988, the concern of incorrectly measuring entry because it happened before the beginning of the time span of the data is limited.

⁶87% of workers holding a higher-level occupation at entry holds a higher-level occupation five years later.

from OECD (2002). Table 2 reports the list of ICT industries together with their share in total employment and in skilled employment over the sample period.

The ICT sector is intensive in skilled labor. This can be seen in Table 2 which shows that the ICT sector represents 5.4% of aggregate employment and 14.9% of skilled employment during the sample period. The share of skilled workers is 15% over all industries vs. 42% in the ICT sector. In terms of education attainment, for which we have information only for a subset of individuals as explained in Section 2.1, the fraction of workers holding a master’s degree is 14% over all industries vs. 28% in the ICT sector.

The ICT sector is more specifically intensive in science/engineering skills. The fraction of skilled workers holding an occupation intensive in science/engineering skills is 35% across all industries vs. 70% in the ICT sector.

3.2 The ICT Boom and Bust

Figure 1 illustrates the boom and bust cycle in the ICT sector in the late 1990s. Panel A of the figure plots the overall employment share of the ICT sector. It increases from approximately 5% in 1995 up to 6% in 2001 and back to 5% in 2004.

Panel B shows that, as expected given the over-representation of skilled workers in the ICT sector, the boom is more pronounced for skilled workers. The share of the ICT sector in total skilled employment goes from 13% in 1996 up to 17% 2001 and down to 15% in 2004. The overall pattern of the ICT sector’s share in skilled employment is a long-term upward trend with a sharp upward deviation from the trend during the 1998–2001 period: the ICT boom.

Panel C shows that the deviation from the trend in the ICT sector’s share of skilled employment is 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 those who have been in the labor force for five years or more.⁷ While the latter exhibits an upward trend, it shows no significant deviation from the trend during the ICT boom. In contrast, the component representing workers recently entered into the labor force exhibits a sharp upward deviation from the trend during the ICT boom. This pattern is naturally explained by the notion that individuals can decide in which sector to start their career but inter-sector mobility after entry is not as easy.

Panel D plots the share of skilled labor market entrants starting in the ICT sector. Again consistent with sectoral choices being more responsive to market conditions at the time of labor market entry, the share of skilled entrants starting in the ICT sector exhibits a much sharper deviation from the trend during the ICT boom than that for the entire population of skilled workers in Panel B. The fraction of skilled entrants starting in the ICT sector almost doubles between 1996

⁷We reiterate that, as in the rest of the paper, entry in the labor force is defined as the individual taking her first full-time job lasting for at least six months.

and 1999, from 18% to 33%, before dropping down to 20% in 2004.

Three conclusions emerge from Figure 1. First, there is a boom-bust cycle in the ICT sector, characterized by significant changes in labor allocation towards and then away from the ICT sector. Second, these allocation changes are concentrated on skilled labor market entrants. Third, the time span of the ICT sector boom is quite sharply delimited with a boom phase starting in 1997-98 and a bust in 2002. This feature of the ICT boom allows us to define an “ICT boom cohort” of skilled workers, whose long-run wage dynamics we study in the next section.

4 Wage Dynamics of the Boom Cohort

4.1 Panel Analysis

We study the wage dynamics of skilled workers who enter the labor market during the ICT boom. We define the boom cohort as the set of skilled workers entering the labor market during the years 1998 to 2001. To start with, we split the boom cohort into the early boom cohort (1998–99) and late boom cohort (2000–01) and estimate the following wage equation separately for each of them:

$$\log(w_{i,t}) = \alpha_t + \beta_t ICT_{i,0} + \gamma_t X_i + \epsilon_{i,t}, \quad (1)$$

where $w_{i,t}$ is the annualized wage of worker i in year t , $ICT_{i,0}$ is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector, and X_i collects a set of worker characteristics: sex, age and age squared at entry, entry year, and two-digit occupation at entry. β_t measures the wage differential in year t for an individual who started in the ICT sector relative to an individual of the same cohort and with the same observable characteristics who started outside the ICT sector.

Figure 2 plots the time-series of β_t for the boom cohorts. The two boom cohorts 1998–99 (blue line) and 2000–01 (red line) exhibit similar patterns. Workers starting in the ICT sector during the early phase of the boom earn an entry wage on average 5% higher than workers of the same cohort and with the same observable characteristics starting outside of the ICT sector. The wage difference vanishes rapidly as the boom ends (and maybe even before the bust). Perhaps more surprisingly, this wage difference keeps falling after the bust and turns negative. By 2010, workers who started in the ICT sector have a wage discount of 5% on average relative to workers of the same cohort who started outside the ICT sector. The pattern is similar for the late boom cohort.

Since the two cohorts have similar wage patterns, Table 3 report regression results pooling all entrants during the ICT boom period 1998–2001. We estimate equation (1) using, for each worker, the year of entry and the years 2002, 2006 and 2010. Accordingly, the regression equation includes four coefficients β_t measuring the wage difference between workers who started in the ICT sector and those who did not, at the time of entry, in 2002, in 2006, and in 2010. Column (1) shows that

during the boom, entrants in the ICT sector start with an average higher wage of 3.7% (significant at 1%) relative to entrants in other sectors. This higher wage decreases and becomes negative over time. In 2010, these workers earn on average 5.6% less (significant at 1%) than workers who started outside of the ICT sectors.

When an individual has no full-time employment spell lasting for at least six months in a given year, the associated worker-year observation is missing, making the panel unbalanced. It implies that the pool of workers on which β_t is identified changes through time. Thus, the level of β_t may not be comparable across years if the propensity to drop from full-time employment is different among workers who started in the ICT sector and those who did not.

In column (2), we control for such composition effects by adding worker fixed effects in equation (1). In this case, the β_t time-series is only identified up to an additive constant. We take the entry year as the baseline and estimate the wage difference between entrants in the ICT sector and other entrants relative to the wage difference at entry. The result is similar: the wage difference decreases over time after the ICT sector bust and reaches -8.9% (significant at 1%) in 2010. The wage discount in 2010 is larger with worker fixed effects because it is defined relative to the wage difference at the time of entry during the boom, which is positive.

4.2 Long-Difference Analysis

Since individuals entering into the ICT sector during the boom experience a steady wage decline after the bust relative to same-cohort individuals starting in other sectors, we now focus on the long difference in log wage between the entry year and 2010. We estimate the cross-sectional regression

$$\log(w_{i,2010}) - \log(w_{i,0}) = \beta ICT_{i,0} + \gamma X_i + \epsilon_i. \quad (2)$$

The identification of β in equation (2) comes from the same variation in the data as the identification of β_{2010} in equation (1) with worker fixed effects and taking the year entry as the reference year. The coefficient on ICT_0 in column (1) of Table 4 implies that entrants in the booming ICT sector experienced lower wage growth from entry to 2010 by 8.5% (significant at 1%).⁸

We allow for regional disparities in wage dynamics in column (2) and include *department* fixed effects.⁹ The baseline result remains and is even slightly stronger, reflecting the facts that the ICT sector is over-represented in urban areas and that wage growth has been stronger in urban areas in the last decades. In the rest of the table, we address three concerns. First, it could be that the negative long-term wage growth of ICT entrants relative to entrants in other sectors is driven by

⁸The coefficient is not exactly equal to the one on $ICT_0 \times (t = 2010)$ in column (2) of Table 3 because the latter depends on worker fixed effects that are estimated using the years of entry, 2002, 2006 and 2010, whereas the coefficient in column (1) of Table 4 is estimated only using the years of entry and 2010.

⁹Departments are administrative entities that split the French territory into 99 distinct geographical zones. The average population of a department is 600,000 inhabitants.

very high wage growth in a few other sectors, such as the finance sector (see Philippon and Reshef (2012) for the US and Célérier and Vallée (2017) for France). To rule out this concern, we exclude entrants starting in the finance sector (they represent 5% of skilled entrants). The relative negative long-term wage growth of ICT entrants is slightly reduced (column (3)), reflecting the high wage growth in finance during the 2000s, but it remains large at -7.7% and statistically significant at 1%. A second concern could be that the characteristics of workers entering in the ICT sector during the boom worsen. To test whether it is the case, we exploit the 4/30th subsample of the panel data that is matched with Census information. We retrieve the level of education and construct two variables of educational attainment: a dummy variable equal to one if the individual holds at least a three-year college degree (*Licence* or equivalent) and a dummy variable equal to one if the individual holds at least a five-year college degree (*Master* or equivalent). We do not consider whether the individual graduated from high school because almost all skilled workers entering the labor force in this period do. 91% of skilled entrants hold at least a three-year college degree and 83% hold at least a five-year college degree. We first re-estimate equation (2) on the small subsample of entrants matched with Census data. The ICT discount in column (4) remains significant at 1% and is slightly larger than that on the main sample due to sampling errors (the p -value for the difference between the coefficient estimated on the full sample and the coefficient estimated on the smaller sample matched with Census data is 0.2). In column (5), we control for the level of education and it does not change the result.

A third concern might be that the income of entrants who become entrepreneurs may be underestimated because the employer-employee data do not include capital income. This measurement error can bias downwards the estimate of β in equation (2) if individuals starting in the ICT sector during the boom are more likely to become entrepreneur than entrants in other sectors. To test whether this is the case, we construct a dummy variable equal to one if the individual is an entrepreneur in 2010, defined as the two-digit occupation code 23 “head of company with at least ten employees”.¹⁰ It is the case for 1.2% of skilled entrants during the boom. Column (6) shows that, using the entrepreneur dummy as the dependent variable, the coefficient on ICT_0 is negative and borderline significant. That is, if anything, individuals starting in the ICT sector during the boom are less likely to become entrepreneurs by half a percentage point.¹¹ Overall, the long-run wage discount experienced by entrants in the booming ICT sector is not explained by mismeasurement of entrepreneurial income.

¹⁰We obtain a similar result when using the broader one-digit occupation 3 “head of company of less or more than 10 employees” (untabulated).

¹¹One caveat is that if an entrepreneur does not pay herself a wage (only dividends), she does not appear in the employer-employee data. However, if entrepreneurs’ propensity to pay themselves a wage is not correlated with her entering into the ICT sector, then our conclusion that individuals starting in the ICT sector are not more likely to become entrepreneurs is robust to this limitation of the data.

4.3 Quantiles Analysis

So far, we have shown that a career start in the ICT sector during the boom is associated with a robust negative relative wage growth. One possible interpretation of this result is that it was risky to start in the booming ICT sector given the uncertainty regarding which firms and technologies will prevail in the long run. In this case, akin to patterns documented in the literature on the returns to entrepreneurship (e.g. Hamilton (2000), Hurst and Pugsley (2015)), the negative mean in relative wage growth may conceal a more nuanced pattern across the wage growth distribution. In particular, the low mean may be associated with a low probability of success, negative skewness, and positive relative wage growth in the right tail of the distribution.

Table 5 reports estimates of quantile regressions for the 10th, 25th, 50th, 75th, and 90th percentiles of wage growth, including the same set of control variables as in equation (2). The main pattern that emerges is that workers starting in the ICT sector during the boom experience robust negative wage growth across the whole wage growth distribution, with long-run discounts ranging from 7% (for the 50th percentile) to 11% (for the 90th percentile). If anything, the wage growth discount is larger at the top of the wage growth distribution, clearly rejecting the hypothesis that the low mean wage growth is associated with a small probability of very positive outcomes. Thus, the boom in the ICT sector does not appear to be a period creating winners and losers among skilled labor market entrants (possibly with an over-presentation of losers). Instead, it has shifted the whole long-run wage distribution to the left for talents who allocated themselves into the booming ICT sector.

The results of the quantile regressions have another important implication. They suggest that the long-run wage growth discount associated with starting in the ICT sector we have uncovered is not explained by the fact that the marginal worker who selects into the booming ICT sector is of worse (unobservable) quality. Suppose it was the case and the pool of workers who selected into the ICT sector during the boom consists of the set of workers who would have gone into the ICT sector no matter what and a set of low-quality workers who self-select into the ICT sector because of the boom. This shift in the worker quality distribution would add a mass to the left of the wage growth distribution, shifting the bottom quantiles to the left by much more than the top quantiles. This hypothesis is clearly rejected by the results of the quantile regressions showing that, if anything, the top quantiles drop by more. In the next section, we provide another piece of evidence that is also inconsistent with a selection effect.

5 A Treatment Effect of the Boom

In this section, we study whether the long-run wage growth discount for entrants in the booming ICT sector documented in Section 4 reflects a treatment effect of the technological boom and bust this sector experienced. To this aim, we investigate two alternative explanations. In Section 5.1, we

study whether this pattern can be explained by a composition effect by which the marginal skilled worker attracted by the booming sector has a low intrinsic productivity. In Section 5.2, we analyze whether the ICT sector experiences a cycle during which the baseline productivity in the sector is low overall but is temporarily high during the boom period.

5.1 Ruling Out a Selection Effect

The negative relative wage growth for entrants who started in the ICT sector during the boom may be explained by the fact that the marginal skilled worker attracted by the booming sector has a low intrinsic productivity. The low relative wage these workers should earn would be masked during the boom period, but as the boom ends and wages start reflecting workers' productivity more accurately, the higher relative wage for entrants in the ICT sector would then turn into negative relative wage over time. In other words, the long-term ICT wage discount may reflect a selection effect rather than the treatment effect of entering into the booming ICT sector.

To test this hypothesis, we exploit the fact that the time span of the ICT boom is well delimited. Thus, individuals entering in the labor market in the period preceding the boom are unlikely to select into the ICT sector because of the boom. Yet, this “pre-boom” cohort who started in the ICT sector will experience the boom. The comparison of the long-term wage dynamics of individuals starting in the ICT sector just before the boom and those starting in the ICT sector during the boom will thus allow us to disentangle between a treatment effect and a selection effect. Similar long-term wage dynamics of the pre-boom cohort and of the boom cohort would be consistent with a treatment effect of the boom. In contrast, different long-term dynamics for the pre-boom and boom cohorts would be consistent with a selection effect.

Figure 3 shows the wage dynamics of the pre-boom cohorts 1994–95 and 1996–97. We estimate equation (1) for each cohort and plot the time-series of β_t . The figure shows that workers entering in the ICT sector during the mid-1990s, i.e., before the ICT boom, earn salaries similar to that of workers entering in other sectors until the beginning of the boom. Then, they experience fast wage growth during the ICT boom. The ICT-sector premium reaches 5% at the peak of the boom, similar to that of the boom cohorts in Figure 2. Crucially, when the boom ends, pre-boom cohorts experience a similar wage dynamics to that of boom cohorts: wages relative to entrants in other sectors keep falling over time and end up negative. By 2010, workers who started in the ICT sector before the boom have a wage discount of 7% on average relative to workers of the same cohort who started outside the ICT sector.

Table 6 mirrors Table 3 for the pre-boom cohort and reports the estimated β_t for the year of entry and the years 1998, 2000, 2002, 2006 and 2010. Since the 1994–95 and 1996–97 cohorts exhibit similar patterns in Figure 3, we pool them together in the regression analysis. We also exclude the 1997 cohort because it might be argued that the ICT boom has already started in 1997 (see Figure 1)

and define the pre-boom cohort as workers starting during 1994–1996.¹² Column (1) shows that the wage dynamics observed in Figure 3 reflects statistically significant effects. Consistent with the idea that the entrants who started in the ICT sector before the boom have similar productivity, their relative wage is similar to those who started in other sectors until 1997. Then, at the pick of the boom in 2000, entrants in the ICT sector have an average wage higher by 4.8% (significant at the 1% level). The relative wage of entrants in the ICT sector then declines and becomes negative over time. In 2010, these workers earn on average 8.5% less (significant at the 1% level) than workers who started outside of the ICT sector.

Worker fixed effects are included in the specification of column (2) to control for potential composition effects as discussed in Section 4.2. The entry year is used as the baseline to estimate wage differences between entrants in the ICT sector and other entrants. The pattern is robust. We observe a large positive wage difference of 5.9% (significant at 1%) in 2000. The wage difference then decreases after the bust and reaches -7.2% (significant at 1%) in 2010.

Overall, the evidence is not consistent by which a selection effect where the booming ICT sector selects individuals with lower intrinsic productivity. Instead, the evidence is more consistent with a treatment effect on workers experiencing the ICT boom.

5.2 Ruling out a Cycle Effect

Another explanation for our result could be that the baseline productivity in the ICT sector is low overall, but this low productivity would be masked by the boom in the late 1990s. When the boom ends, the wage of workers who started in the ICT sector would go back to reflecting the low productivity of the sector and end up below that of workers who started in other sectors. Note first that this explanation is not entirely consistent with the fact that the entry wage of entrants in the ICT sector during the pre-boom period does not display a discount relative to entrants in other sectors, suggesting that the ICT sector did not have a lower productivity before the boom. Yet, it might still be the case that productivity in the ICT sector experiences an overall decline after the bust in 2002.

To test for such a cycle effect, we now analyze the post-boom cohorts 2002–03 and 2004–05. If the ICT sector experiences a decrease in productivity after the bust, post-boom cohorts should also exhibit a decline in their relative wages.

Figure 4 shows the relative wage dynamics of the post-boom cohorts, which we obtain by estimating (1) on these cohorts and plotting the time-series of β_t . While entrants in the ICT sector from the 2002–03 cohort, and even more for the 2004–05 cohort, start with lower wages, the wage discount progressively vanishes such that, by 2010, workers who started in the ICT sector in the post-boom period earn approximately the same (for the 2004–05) as or even more (for the 2002–03

¹²The results in this section are robust to including 1997 in the pre-boom period.

cohort) than workers who started in other sectors.

Table 7 reports the estimated β_t for the year of entry and the years 2006 and 2010. We also pool the two post-cohorts together and exclude 2002 from the post-boom period in order to leave a gap year between the boom period and post-boom period. Accordingly, from now on, we refer to 2003–2005 as the post-boom period.¹³ Column (1) shows that post-boom entrants starting in the ICT sector earn a 4% lower wage (significant at 1%) than entrants in other sectors. Crucially, their relative wage show no further decline and, to the contrary, they recover over time and converge towards the same wage level as in other sectors, such that by 2010, there is no longer any significant difference. Column (2) shows a similar pattern when we include worker fixed effects and estimate the wage dynamics relative to the entry wage. Entrants in the ICT sector experience faster wage growth from entry to 2010 by 2.9% (significant at 5%) relative to entrants in other sectors.

This evidence is inconsistent with a secular decline in productivity in the ICT sector during the 2000s driving the wages of all skilled workers down. Instead, the evidence that the post-boom cohort experiences an opposite wage dynamics to that of the pre-boom and boom cohorts is more consistent with a treatment effect affecting workers exposed to the ICT sector during the boom.

6 Economic Mechanisms

In this section, we hypothesis that the technology boom was characterized by a high rate of experimentation and technological change, causing workers who start in the ICT sector during the boom to accumulate human capital early in their career, that is likely to have low value and to weigh on their productivity in long run. This explanation would be consistent with the wage dynamics we observe. Both the pre-boom and boom cohorts of workers starting in the ICT sector are exposed to this period of radical technological change during which they acquire knowledge that deteriorate their long-run productivity relative to workers starting in other sectors. In contrast, workers starting in the ICT sector after the sector stabilizes are not exposed to this period of low value of on-the-job learning, and thus do not experience any long-run wage decline.

We explore two specific channels by which human capital accumulated during the ICT boom may have low long-term value. First, firms in the ICT sector may be more likely to fail or downsize because of high entry and investment during the boom, thus forcing higher turnover for their employees, and in particular for their recently hired employees.¹⁴ If workers accumulate firm-specific human capital, a higher probability of forced job termination will lead to a loss of human capital and lower subsequent wage growth (see, e.g., Jarosch (2015)). This forced job termination might also increase the risk of mismatched jobs, which could impact human capital and limit possibilities for career advancement.

¹³The results in this section are robust to including 2002 in the post-boom period.

¹⁴Labor regulation in France is such that workers accumulate rights in proportion to their time within the firm. As a result, firing recently hired employees is less costly for firms.

Second, human capital accumulated by workers in the ICT sector during the boom may have low value because technology changes rapidly and skill obsolescence accelerates during technology booms. Individuals starting in the ICT sector during the boom learn about technologies that rapidly become obsolete, reducing their long-term productivity. Under this hypothesis, skill obsolescence is not necessarily related to individuals changing employer and can happen within job spells.

6.1 Within-Jobs vs. Between-Jobs Wage Growth

As a first pass to disentangle between these two channels, we focus again on the boom cohort and decompose workers' wage growth from entry to 2010 into a within-jobs component and a between-jobs component. Indexing by $t = 0, \dots, T$ the years in which we observe worker i in the data and denoting by $E_{i,t}$ her employer in year t , we define within-jobs wage growth as

$$\sum_{t=1}^T 1_{E_{i,t}=E_{i,t-1}} [\log(w_{i,t}) - \log(w_{i,t-1})] \quad (3)$$

and between-jobs wage growth as

$$\sum_{t=1}^T 1_{E_{i,t} \neq E_{i,t-1}} [\log(w_{i,t}) - \log(w_{i,t-1})]. \quad (4)$$

We estimate equation (2) using these two components of wage growth as dependent variables. If entrants in the booming ICT sector experience lower wage growth because they are forced to change jobs and in the process lose human capital, we expect the wage growth discount to be come predominantly from the between-jobs component.

Table 8 shows that the wage growth discount comes in fact almost entirely from the within-jobs component. Of the total 8.5% wage growth discount experienced by entrants in the booming ICT sector, 7.7 percentage points (significant at 1%) come from slower wage growth within job spells and only 0.8 percentage points (insignificant) come from slower wage growth during job transitions. Note that this result does not reflect the fact that wage growth happens almost only within job spells unconditionally. For skilled workers entering the labor market (in any sector) during the boom period, within-jobs and between-jobs wage growth explain respectively 36% and 23% of the variation in total wage growth.¹⁵ This preliminary evidence is thus more consistent with entrants in the booming ICT sector seeing their human capital depreciating over time independently of potential job terminations.

The absence of a negative wage growth at the time of job transitions may nevertheless underestimate the effect of job terminations on long-run wage growth. For instance, job changes may

¹⁵The R2s do not sum to one because within-jobs wage and between-jobs wage growth are negatively correlated in the cross-section of workers.

induce a lower probability of promotion in the long-run even if the job transition themselves are not associated with negative wage changes. Forced terminations may also increase the risk of mismatch in the new job and thus negatively impacting human capital accumulation and limit the possibilities for career advancement. In the next section, we seek to detect more accurately forced job termination and test if it can explain the long-run wage growth discount experienced by entrants in the ICT sector during the boom.

6.2 Job Termination

The boom in the ICT sector in the late 1990s was associated with a rapid growth in investment and hiring, resulting in the bust of 2001–2002. During this period, many workers had to change jobs as firms in the ICT sector shrank as we documented in Section 3.2. This can explain the drop in long-term relative wages if job displacement entails losing human capital that is specific to the firm. In this case, even when the economy or the ICT sector picks up, displaced workers will still earn lower wages as their firm-specific human capital has been lost.

This mechanism would explain the long-run wage discount of the pre-bubble cohorts and the bubble cohorts, as well as the absence of a discount for the post-bubble cohort. It also makes two additional testable predictions. First, workers starting in the ICT sector should be more likely to be forced to switch employers. Second, the wage discount should be associated with forced job termination.

We construct several variables to measure (forced) job termination. The first is a dummy equal to one if the worker switches to another employer within the first four years after entry. It is the case for 56% of skilled entrants over 1994–2005 on average across sectors. Table 9, column (1) shows how the probability of job termination for entrants depends on the sector of entry for the pre-boom, boom, and post-boom cohorts. Specifically, we regress the job termination dummy on ICT_0 interacted with the dummy variables for the pre-boom cohort dummy (entry year 1994–1996), the boom cohort dummy (entry year 1998–2001), and the post-boom cohort dummy (entry year 2003–2005), and the same set of controls as before (sex, age and age squared at entry, and two-digit occupation at entry, all interacted with the cohort dummies, and entry year fixed effects).¹⁶ The first insight from column (1) is that, across all cohorts, skilled workers starting in the ICT sector are always more likely to switch employer within the first four years after entry, whether these job terminations are voluntary or forced. The second insight is that this effect is stronger for the boom cohort. Consistent with the ICT sector experiencing a boom-bust cycle, workers starting in the ICT sector during the boom have a 2.3 percentage points higher probability to switch employer relative to entrants in the ICT sector before or after the boom, and relative to entrants in other sectors, although the difference is not statistically significant (p -value 0.28).

¹⁶The regression does not have non-interacted cohort dummies because they are absorbed by entry year fixed effects.

In column (2) we use a second measure of job termination defined at a longer horizon. We construct a dummy equal to one if the individual has switched employer by 2010. It is the case for 76% of skilled entrants. Note that in this case, the horizon at which job termination is defined varies across cohorts. The result is in line with the result using the first measure of job termination. Workers entering in the ICT sector are more likely to have left their initial employer by 2010, and more so for entrants in the ICT boom period. Workers starting in the ICT sector during the boom have a 5.4 percentage points (significant at 5%) higher probability to have switched employer by 2010 relative the workers starting in the ICT sector before the boom and relative to entrants in other sectors. Relative to the post-boom cohort, the probability is higher by 3.5 percentage points (p -value of 0.15).

An important limitation of our first two measures of job termination is that they do not distinguish between voluntary job termination and forced job termination. To gain more insight into the effect of forced job termination, we construct two proxies for forced job termination. The first one considers that a job termination is forced if it is associated with a transition to a lower-paid job. Then, for each worker, we construct a forced job termination dummy equal to one if she experiences a forced job termination according to this definition within the first four years after entry. It is the case for 15% of skilled entrants over the sample period. We use this forced job termination dummy as the dependent variable in column (3). Interestingly, it shows that the higher probability of job termination for the ICT boom cohort documented in column (1) is (more than) entirely explained by forced terminations. entrants in the ICT sector during the boom are 6.1 percentage points (significant at 1%) more likely to experience a forced job termination in the first years of their career than entrants in other sectors, whereas there is no significant differences for the pre-boom cohort and the post-boom cohort (the difference in coefficient between the boom and pre-boom cohorts is significant at 5% and the one between the boom and post-boom cohorts at 1%). The fact that the difference in forced job termination between the ICT boom cohort and the pre- or post-boom cohorts is larger than that in any job termination suggests that voluntary job termination decreases for workers starting in the ICT sector during the boom. This is consistent with reduced job opportunities during the bust of the ICT sector.

In column (4) we use a second proxy for forced job termination. We consider that a job termination is forced if the employer has negative employment growth in the year in which the job is terminated. Again, we construct a dummy equal to one if the worker experiences a forced job termination according to this definition within the first four years after entry. It is the case for 19% of skilled entrants over the sample period. The result is in line with the one using the first proxy for forced termination. Overall, the pattern is consistent with an increase in forced job termination for workers starting in the ICT sector during the boom.

The next question is whether the higher displacement risk faced by workers starting in the booming ICT sector explains their long-term decline in relative wage. The negative effect of forced

job termination on long-term wages may arise in two cases. First, it may arise independently of the worker’s sector. Second, it may arise more specifically when a worker loses his job in a sector that experiences a bust. We consider these two cases in turn.

Panel A of Table 10 shows that controlling for job termination explains a negligible part of the ICT boom cohort’s long-run wage discount. We focus again on the boom cohort and augment the baseline wage regression (2) by controlling for each of the four measures of job termination. Compared to the baseline long-run discount of 8.5% (column (1) of Table 4), controlling for job termination explains at most 0.9 percentage points of this discount, when using the measure of forced termination associated with a wage in the subsequent job in column (3).¹⁷

Panel B of Table 10 explores the possibility that the effect of job displacement on long-term wage growth of entrants during the boom period is different in the ICT sector and in other sectors. This can happen if job displacement is more detrimental in sectors experiencing a bust. To allow for this possibility, we additionally control for the interaction term ICT_0 times job termination. In this case, the coefficient on (non-interacted) ICT_0 can be interpreted as the wage growth difference between entrants starting in the ICT sector and experiencing no job termination, and entrants in other sectors experiencing no job termination. With all for proxies of job termination, we find a wage discount of the same magnitude (between 7.4% and 10%) as in the baseline specification (8.5%). A particularly telling result is the one reported in column (2), showing that workers starting in the ICT sector during the boom and still working with their initial employer in 2010, experience a 10 percentage points slower wage growth than entrants in other sectors and also working with their initial employer in 2010.

Overall, the evidence is inconsistent with the wage discount of the ICT boom cohort being explained by workers being more likely to lose their job early in their career. Even entrants in the ICT sector who do not have to switch employers experience as poor a long-term wage growth as those losing their jobs. It suggests that these workers experience a decline in their long-run productivity that goes beyond firm-specific effects and that shows up whatever their future career path.

6.3 Skill Obsolescence

We now explore the idea that workers starting in the ICT sector during the boom accumulate human capital early in their career that rapidly becomes obsolete. If the long-run relative wage growth decline of these individuals is explained by this mechanism, we expect it to be stronger for workers holding jobs with a higher technological content or working in firms more intensive in technology, because the human capital accumulated on these jobs depreciates faster as technology changes.

We test this hypothesis using several proxies for jobs’ technological content. The first proxy

¹⁷Including all four proxies of job termination in the same regression gives a similar result, with a coefficient on ICT_0 equal to -0.077 (significant at 1%).

is the occupation held by labor market entrants. The two-digit occupation classification in the employer-employee data distinguishes between occupations with a science/engineering skill content (hereafter “engineers”) and those with a management/business content (hereafter “managers”). We construct a variable $Engineer_i$ equal to one if the occupation of worker i in her first job is as an engineer and equal to zero otherwise. We expect the long-run wage discount to be stronger for workers starting in the ICT sector during the boom as engineers than for managers. The second proxy aims at capturing the technological intensity of firms in which entrants start their career. We define $TechFirm_i$ as the fraction of engineers in the initial employer of worker i . The third proxy aims at capturing the technological intensity of specific (four-digit) sectors of the broad ICT sector in which entrants in the ICT sector start their career. We define $TechSector_i$ as the fraction of engineers in the four-digit sector in which worker i holds her first job.

Table 11 shows how long-run wage growth depends on jobs’ technological content. In column (1), we estimate equation (2) adding the interaction term $ICT_{i,0} \times Engineer_i$ as an explanatory variable.¹⁸ The coefficient on the interaction term is negative and significant at the 5% level, while the coefficient on the non-interacted term ICT_0 is not significant. Thus, the long-run wage discount is concentrated on engineers, while managers experience a much smaller and statistically insignificant relative wage decline. In contrast, engineers who started in the ICT sector experience slower wage growth by 6.6 percentage points relative to managers who started in the same sector.

We use the proxy for the firm’s technological intensity in column (2) and include the interaction of ICT_0 with $TechFirm$ as an explanatory variable. The coefficient on the interaction term is negative and significant at the 5% level. Thus, the long-run wage discount for workers starting in the ICT sector during the boom is stronger for those who started in more-tech firms. To check that the result is not driven by a more general pattern by which workers starting in more-tech firms even outside the ICT sector would experience slower wage growth, we control in column (3) for $TechFirm$ not interacted with ICT_0 . Two results appear. First, the coefficient on the interaction term is barely affected by the inclusion of this control. Second, the coefficient on the non-interacted term $TechFirm$ is not significant. Thus, patterns of wage dynamics reflect the rapid obsolescence of technologies developed specifically in the ICT sector during this period rather than a general trend of technological skills obsolescence.

A similar pattern emerges when we use the proxy for the sector’s technological content. Columns (4) shows that the long-run wage discount for workers starting in the ICT sector during the boom is stronger for those who started in more-tech sectors. Column (5) shows that the result is not explained by the fact that workers starting in more-tech sectors even outside the ICT sector experience slower wage growth.

One final concern might be that engineers in the ICT sectors, or more generally workers in

¹⁸We do not include the non-interacted term $Engineer_i$ because the baseline specification already includes fixed effects for the initial occupation.

technology-intensive ICT firms or sectors, always have slower wage growth for reasons unrelated to the ICT boom. Table 12 replicates the same analysis for the post-boom cohort and shows that it is not the case. Column (1) shows that the wage trajectory of engineers starting in the ICT sector during 2003–2005 is not significantly different from that of managers. Columns (2) and (3) shows that a similar conclusion holds for workers starting in technology-intensive firms. Columns (4) and (5) show that, if anything, workers starting in more technology-intensive ICT sectors during the post-boom period experience faster wage growth than those starting in less technology-intensive ICT sectors.

7 Concluding Remarks

In this paper, we show that entrants who started during the ICT boom of the late 1990s end up in the long run with a 8 percentage points lower wage relative to similar skilled workers who started in a different sector.

We rule out that the effect is driven by marginal entrants in the ICT sector being of lower quality by first showing that this negative relative wage hold with a similar magnitude across the whole wage distribution, suggesting that the average effect we observe is not driven by an increase in the mass of less productive workers. Second, we show that workers who started in the ICT sector during the pre-boom period experience a similar long-run negative relative wage, while displaying similar wages to that of workers in other sector before the boom starts.

We conclude by offering suggestive evidence that entrants in the booming ICT sector were exposed to technologies that became rapidly obsolete. By contrast, the forced job termination some workers experienced when the sector shrank after the boom ends, does not explain the poor long-term wage growth of entrants in the booming ICT face.


Overall, the results suggest that boom-burst episodes can have large and long-lasting effects on skilled workers' accumulation of human capital, even when they happen in technology-intensive sectors.

References

- Altonji, Joseph G, Lisa B Kahn, and Jamin D Speer. 2016. “Cashier or consultant? Entry labor market conditions, field of study, and career success.” Journal of Labor Economics 34 (S1):S361–S401.
- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. “The wage policy of a firm.” Quarterly Journal of Economics 109 (4):921–955.
- Beaudry, Paul and John DiNardo. 1991. “The effect of implicit contracts on the movement of wages over the business cycle: evidence from micro data.” Journal of Political Economy 99 (3):665–688.
- Beaudry, Paul, David A Green, and Benjamin M Sand. 2016. “The great reversal in the demand for skill and cognitive tasks.” Journal of Labor Economics 34 (S1):S199–S247.
- Brown, J. R., S. M. Fazzari, and B. C. Petersen. 2009. “Financing innovation and growth: cash flow, external equity, and the 1990s R&D Boom.” Journal of Finance 64 (1):151–185.
- Campello, Murillo and John R Graham. 2013. “Do stock prices influence corporate decisions? Evidence from the technology bubble.” Journal of Financial Economics 107 (1):89–110.
- C  lerier, Claire and Boris Vall  e. 2017. “Returns to talent and the finance wage premium.” Working Paper .
- Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo. 2015. “Housing booms and busts, labor market opportunities, and college attendance.” Tech. rep., National Bureau of Economic Research.
- Choi, Darwin, Dong Lou, and Abhiroop Mukherjee. 2017. “The effect of superstar firms on college major choice.” .
- Devereux, Paul. 2002. “Occupational upgrading and the business cycle.” Labour 16 (3):423–452.
- Dong, Ming, David Hirshleifer, and Siew Hong Teoh. 2017. “Stock Market Overvaluation, Moon Shots, and Corporate Innovation.” NBER Working Paper .
- Gibbons, Robert and Michael Waldman. 2006. “Enriching a theory of wage and promotion dynamics inside firms.” Journal of Labor Economics 24 (1):59–107.
- Hamilton, Barton H. 2000. “Does entrepreneurship pay? An empirical analysis of the returns to self-employment.” Journal of Political economy 108 (3):604–631.
- Hurst, Erik G. and Benjamin W. Pugsley. 2015. “Wealth, Tastes, and Entrepreneurial Choice.” NBER Working Paper .


- Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993. "Earnings losses of displaced workers." The American economic review :685–709.
- Jarosch, Gregor. 2015. "Searching for Job Security and the Consequences of Job Loss." Working Paper .
- Kahn, Lisa B. 2010. "The long-term labor market consequences of graduating from college in a bad economy." Labour Economics 17 (2):303–316.
- Kerr, William R, Ramana Nanda, and Matthew Rhodes-Kropf. 2014. "Entrepreneurship as experimentation." Journal of Economic Perspectives 28 (3):25–48.
- Manso, Gustavo, Benjamin Balsmeier, and Lee Fleming. 2017. "Heterogeneous Innovation over the Business Cycle." Working Paper .
- Nanda, Ramana and Matthew Rhodes-Kropf. 2013. "Investment cycles and startup innovation." Journal of Financial Economics 110 (2):403–418.
- OECD. 2002. Measuring the Information Economy. OECD Publications.
- Ofek, Eli and Matthew Richardson. 2003. "Dotcom mania: The rise and fall of internet stock prices." The Journal of Finance 58 (3):1113–1137.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. "Short and long-term career effects of graduating in a recession." American Economic Journal: Applied Economics 4 (1):1–29.
- Ouimet, Paige and Rebecca Zarutskie. 2014. "Who works for startups? The relation between firm age, employee age, and growth." Journal of financial Economics 112 (3):386–407.
- Philippon, Thomas and Ariell Reshef. 2012. "Wages and human capital in the US finance industry: 1909–2006." The Quarterly Journal of Economics 127 (4):1551–1609.
- Schumpeter, Joseph A. 1942. Capitalism, socialism and democracy. Harper & Brothers.
- Topel, Robert H and Michael P Ward. 1992. "Job mobility and the careers of young men." The Quarterly Journal of Economics 107 (2):439–479.
- Violante, Giovanni L. 2002. "Technological acceleration, skill transferability, and the rise in residual inequality." The Quarterly Journal of Economics 117 (1):297–338.

Figure 1: Employment Share of the ICT Sector



Source: Matched employer-employee panel. Panel A plots the share of the ICT sector, as defined by OECD (2002), in total full-time employment. Panel B plots the ICT sector's share in skilled employment, as defined in Section 2.1.

Panel C: Skilled workers: decomposition recent entrants vs. older workers




Panel D: Skilled entrants




Source: Matched employer-employee panel. Panel C decomposes skilled employment in the ICT sector into workers who entered the labor market five years ago or more (blue line) and those who entered four years ago or less (red line). Panel D plots the share of skilled labor market entrants starting in the ICT sector.

Figure 2: Wage Dynamics of ICT Boom Cohorts



The figure displays the β_t coefficient of the wage regression $\log(w_{i,t}) = \alpha_t + \beta_t ICT_{i,0} + \gamma_t X_i + \epsilon_{i,t}$ where $ICT_{i,0}$ is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector and X_i collects control variables listed in Section 4.1. The 1998–99 and 2000–01 cohorts are the skilled workers whose first full-time job was in 1998–99 and 2000–01, respectively.

Figure 3: Wage Dynamics of Pre-Boom Cohorts



The figure displays the β_t coefficient of the wage regression $\log(w_{i,t}) = \alpha_t + \beta_t ICT_{i,0} + \gamma_t X_i + \epsilon_{i,t}$ where $ICT_{i,0}$ is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector and X_i collects control variables listed in Section 4.1. The 1994–95 and 1996–97 cohorts are the skilled workers whose first full-time job was in 1994–95 and 1996–97, respectively.

Figure 4: Wage Dynamics of Post-Boom Cohorts



The figure displays the β_t coefficient of the wage regression $\log(w_{i,t}) = \alpha_t + \beta_t ICT_{i,0} + \gamma_t X_i + \epsilon_{i,t}$ where $ICT_{i,0}$ is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector and X_i collects control variables listed in Section 4.1. The 2002-03 and 2004-05 cohorts are the skilled workers whose first full-time job was in 2002-03 and 2004-05, respectively.

Table 1: Summary Statistics

Panel A: All skilled workers

	N	Mean	P25	P50	P75
Annual wage	1,217,837	41,135	26,136	33,748	46,261
Male	1,217,837	0.72	0	1	1
Age	1,217,837	43	34	42	51

Panel B: Skilled workers entering the labor force over 1994–2005

	N	Mean	P25	P50	P75
Annual wage	160,448	32,512	22,491	28,303	36,987
Male	160,448	0.68	0	1	1
Age at entry	160,448	26	25	26	27

Panel A presents summary statistics at the worker-year level for the period 1994–2010 for the sample of skilled workers in the matched employer-employee panel who hold a full-time job. Panel B reports summary statistics for the subsample of skilled workers who enter the labor force between 1994 and 2005.

Table 2: ICT Industries

ICT industries	ISIC rev 3.1 codes	Share of total employment (%)	Share of skilled employment (%)
ICT: Services		1.9	7.8
IT consultancy	<i>7210</i>	0.7	3.4
Software	<i>7220</i>	0.7	3.2
Data processing	<i>7230</i>	0.3	0.8
Maintenance computers	<i>7250</i>	0.1	0.2
Other data/computer-related services	<i>7123,7240,7290</i>	0.1	0.2
ICT: Telecommunications		1.4	2.2
Telecommunications	<i>6420</i>	1.4	2.2
ICT: Manufacturing		1.7	3.7
Electronic/communication equipment	<i>3210,3220,3230</i>	0.8	1.8
Measurement/navigation equipment	<i>3312,3313</i>	0.5	1.2
Accounting/computing equipment	<i>3000</i>	0.2	0.7
Insulated wire and cable	<i>3130</i>	0.1	0.1
ICT: Wholesale		0.5	1.2
Computers, electronics, telecoms	<i>5151,5152</i>	0.5	1.2
ICT: Total		5.4	14.9

List of ICT industries from OECD (2002). The third (fourth) column reports the 1993–2008 average share in total employment (in skilled employment) of each ICT industry.

Table 3: Wage Dynamics of the Boom Cohort

The table presents the OLS estimates of β_t in equation (1) for skilled entrants of the boom cohort 1998–2001. The dependent variable is log wage of worker i in year t . ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector and $(t=YYYY)$ is a dummy equal one for the year $YYYY =$ entry year, 2002, 2006, and 2010. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. Column (2) includes worker fixed effects and use the year of entry as the baseline year. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage	
	(1)	(2)
$ICT^0 \times (t=0)$.037*** (.0066)	
$ICT^0 \times (t=2002)$.02** (.0082)	-.015** (.0074)
$ICT^0 \times (t=2006)$	-.019* (.011)	-.064*** (.01)
$ICT^0 \times (t=2010)$	-.055*** (.012)	-.089*** (.011)
Worker FE	N	Y
Adjusted-R2	.46	.75
Observations	25,425	23,965

Table 4: Wage Growth of the Boom Cohort

The table presents OLS estimations of equation (2) for skilled entrants of the boom cohort 1998–2001. The dependent variable is wage growth of worker i from entry year to 2010. ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. Column (1) shows the baseline result. Column (2) adds department (administrative geographic areas) fixed effects. Column (3) removes entrants who started in the finance sector. Column (4) restricts the sample to workers we can match with census data. Column (5) includes dummies for degree held by workers. In Column (6), the dependent variable is a dummy equal to one if the worker is an entrepreneur in 2010. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2010 – log wage entry					Entrepreneur
	(1)	(2)	(3)	(4)	(5)	(6)
ICT_0	-.085*** (.012)	-.1*** (.013)	-.077*** (.012)	-.13*** (.035)	-.13*** (.035)	-.0056* (.0034)
Adjusted-R2	.11	.12	.11	.11	.13	.015
Department FE	N	Y	N	N	N	N
Observations	5,097	5,095	4,834	635	635	5,097

Table 5: Quantiles of Wage Growth

The table presents quantile regressions of equation (2) for skilled entrants of the boom cohort 1998–2001. The dependent variable from column (1) to (5) is the 10th, 25th, 50th, 75th, and 90th percentile, respectively, of wage growth of worker i from entry year to 2010. ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Wage growth quantiles				
	P10 (1)	P25 (2)	P50 (3)	P75 (4)	P90 (5)
ICT_0	-.086*** (.023)	-.079*** (.015)	-.07*** (.013)	-.1*** (.014)	-.11*** (.023)
Pseudo-R2	.032	.05	.069	.089	.094
Observations	5,097	5,097	5,097	5,097	5,097

Table 6: Wage Dynamics of the Pre-Boom Cohort

The table presents the OLS estimates of β_t in equation (1) for skilled entrants of the pre-boom cohort 1994–1996. The dependent variable is log wage of worker i in year t . ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector and $(t=YYYY)$ is a dummy equal one for the year $YYYY =$ entry year, 1997, 2000, 2002, 2006, and 2010. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. Column (2) includes worker fixed effects and use the year of entry as the baseline year. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage	
	(1)	(2)
$ICT_0 \times (t=0)$.015 (.011)	
$ICT_0 \times (t=1997)$	-.0011 (.013)	.0026 (.01)
$ICT_0 \times (t=2000)$.048*** (.017)	.059*** (.015)
$ICT_0 \times (t=2002)$.038** (.018)	.026 (.016)
$ICT_0 \times (t=2006)$	-.035* (.02)	-.045** (.018)
$ICT_0 \times (t=2010)$	-.085*** (.024)	-.072*** (.022)
Adjusted-R2	.51	.79
Observations	17,187	16,649

Table 7: Wage Dynamics of the Post-Boom Cohort

The table presents the OLS estimates of β_t in equation (1) for skilled entrants of the boom cohort 2003–2005. The dependent variable is log wage of worker i in year t . ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector and $(t=YYYY)$ is a dummy equal one for the year $YYYY =$ entry year, 2006, and 2010. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. Column (2) includes worker fixed effects and use the year of entry as the baseline year. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage	
	(1)	(2)
$ICT_0 \times (t=0)$	-.04*** (.0096)	
$ICT_0 \times (t=2006)$	-.021* (.012)	.0059 (.009)
$ICT_0 \times (t=2010)$	-.011 (.015)	.029** (.013)
Adjusted-R2	.33	.72
Observations	12,182	11,309

Table 8: Within-Jobs vs. Between-Jobs Wage Growth Decomposition

The table presents the decomposition of workers' wage growth from entry to 2010 into a within-job and a between-job component as defined in equations (3) and (4), for skilled entrants of the boom cohort 1998–2001. ICT_0 is a dummy variable equal to one if i 's first employment spell is in a firm in the ICT sector. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Within-jobs wage growth (1)	Between-jobs wage growth (2)
ICT_0	-.077*** (.012)	-.0076 (.011)
Adjusted-R2	.065	.025
Observations	5,097	5,097

Table 9: Job Termination

The table presents OLS estimations of equation (2) for skilled entrants between 1994 and 2005. ICT_0 is a dummy equal to one if i 's first employment spell is in a firm in the ICT sector. *Pre-boom cohort* is a dummy equal to one if the entry year is 1994–1996, *Boom cohort* equal one if the entry year is 1998–2001, and *Post-boom cohort* equal one if the entry year is 2003–2005. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. The dependent variables are four different proxies for worker job termination. In column (1) it is a dummy equal to one if the worker switches job within the first four years after entry. In column (2) it is a dummy equal to one if the worker has a different employer in 2010 than at entry. In columns (3) and (4), it is a dummy equal to one if the worker switches job during the first four years after entry and this switch is associated with a lower wage (column (3)) or the initial employer has negative employment growth in the year in which the switch happens (column (4)). ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Job termination			
	(1) Within four years	(2) Diff. employer in 2010	(3) Within 4y & Δ wage<0	(4) Within 4y & Δ emp<0
Pre-boom cohort x ICT_0	.087*** (.026)	.038** (.019)	.014 (.018)	-.004 (.022)
Boom cohort x ICT_0	.11*** (.016)	.092*** (.012)	.061*** (.012)	.032** (.014)
Post-boom cohort x ICT_0	.087*** (.023)	.057*** (.021)	-.012 (.017)	.022 (.02)
Adjusted-R2	.022	.059	.0061	.027
Observations	10,602	10,602	10,602	10,602

Table 10: Wage Growth and Job Termination

The table presents OLS estimations of equation (2) for skilled entrants of the boom cohort 1998–2001. The dependent variable is wage growth of worker i from entry year to 2010. ICT_0 is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector. In Panel A, columns (1) to (4), we include the four proxies for job termination used in Table 9 as explanatory variables. In Panel B, columns (1) to (4), we interact each proxy for job termination with ICT_0 to allow the effect of job termination to vary across sectors. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Controlling for job termination

<i>Dependent variable: Wage growth</i>				
<i>Proxy for job termination:</i>	(1)	(2)	(3)	(4)
	Within four years	Diff. employer in 2010	Within 4y & $\Delta\text{wage} < 0$	Within 4y & $\Delta\text{emp} < 0$
ICT_0	-.082*** (.012)	-.083*** (.012)	-.076*** (.012)	-.084*** (.012)
Job termination	-.021* (.011)	-.019 (.013)	-.14*** (.016)	-.013 (.014)
Adjusted-R2	.11	.11	.13	.11
Observations	5,097	5,097	5,097	5,097

Panel B: Controlling for job termination in ICT

<i>Dependent variable: Wage growth</i>				
<i>Proxy for job termination:</i>	(1)	(2)	(3)	(4)
	Within four years	Diff. employer in 2010	Within 4y & $\Delta\text{wage} < 0$	Within 4y & $\Delta\text{emp} < 0$
ICT_0	-.094*** (.017)	-.1*** (.025)	-.074*** (.013)	-.09*** (.013)
Job termination	-.027** (.013)	-.024 (.015)	-.14*** (.022)	-.023 (.019)
ICT_0 x Job termination	.019 (.022)	.024 (.027)	-.012 (.03)	.027 (.028)
Adjusted-R2	.11	.11	.13	.11
Observations	5,097	5,097	5,097	5,097

Table 11: Wage Growth and Job Skill Content

The table presents OLS estimations of equation (2) for skilled entrants of the boom cohort 1998–2001. The dependent variable is wage growth of worker i from entry year to 2010. ICT_0 is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector. We interact ICT_0 with different proxies for the skill intensity of worker i 's job. $Engineer$ equals one if the worker is in an occupation with a science/engineer content (as opposed to management/business). $TechFirm$ is the fraction of engineers in worker i 's initial employer. $TechSector$ is the fraction of engineers in worker i 's initial four-digit sector. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Wage growth				
	(1)	(2)	(3)	(4)	(5)
ICT_0	-.034 (.029)	-.039 (.026)	-.044* (.027)	-.0087 (.031)	-.017 (.032)
$ICT_0 \times Engineer$	-.066** (.032)				
$ICT_0 \times TechFirm$		-.086** (.036)	-.09** (.036)		
$(1-ICT_0) \times TechFirm$			-.029 (.031)		
$ICT_0 \times TechSector$				-.18*** (.065)	-.19*** (.065)
$(1-ICT_0) \times TechSector$					-.081 (.072)
Adjusted-R2	.11	.11	.11	.11	.11
Observations	5,097	4,959	4,959	5,095	5,095

Table 12: Wage Growth and Job Skill Content: Post-Boom Cohorts

The table presents OLS estimations of equation (2) for skilled entrants of the post-boom cohort 2003–2005. The dependent variable is wage growth of worker i from entry year to 2010. ICT_0 is a dummy variable equal to one if worker i 's first employment spell is in a firm in the ICT sector. We interact ICT_0 with different proxies for the skill intensity of worker i 's job. *Engineer* equals one if the worker is in an occupation with a science/engineer content (as opposed to management/business). *TechFirm* is the fraction of engineers in worker i 's initial employer. *TechSector* is the fraction of engineers in worker i 's initial four-digit sector. All regressions include as worker-level control variables: sex, age and age square at entry, two-digit occupation at entry, all interacted with entry year dummy variables, and non-interacted entry year dummy variables. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Wage growth				
	(1)	(2)	(3)	(4)	(5)
ICT_0	.03 (.035)	.009 (.03)	.012 (.03)	-.035 (.04)	-.025 (.041)
$ICT_0 \times \text{Engineer}$.0023 (.038)				
$ICT_0 \times \text{TechFirm}$.052 (.04)	.055 (.04)		
$(1-ICT_0) \times \text{TechFirm}$.017 (.029)		
$ICT_0 \times \text{TechSector}$.14* (.075)	.14* (.075)
$(1-ICT_0) \times \text{TechSector}$.089 (.065)
Adjusted-R2	.093	.098	.098	.095	.095
Observations	3,177	3,066	3,066	3,169	3,169