

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

by

Johan Hombert, HEC Paris and CEPR
Adrien Matray, Princeton University

Griswold Center for Economic Policy Studies
Working Paper No. 255, June 2019

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

Johan Hombert[†] Adrien Matray[‡]

February 2019

Abstract

Using matched employer-employee data from France, we uncover an *ICT boom-cohort discount* on the long-term earnings of the large cohort of skilled workers entering in the Information and Communication Technology (ICT) sector during the 1990s Tech Bubble. Despite starting with 5% higher wages, these workers experience lower wage growth and end up with 6% lower wages fifteen years out, relative to similar workers who started outside the ICT sector. Other moments of the wage distribution are inconsistent with selection effects. We provide suggestive evidence that these workers accumulate human capital early in their career that rapidly depreciates.

*We thank Laurent Bach, Nick Bloom, Denis Gromb, Stefano Lovo, Stefan Lewellen as well as seminar participants at HEC-Paris, BI-Bergen, ESSEC, NOVA Business School and conference participants at the 14th Csef-Igier Symposium on Economics and Institutions, the American Finance Association 2018, AFFI 2018. 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 Julis-Rabinowitz Center for Public Policy and Finance.

[†]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 levels of investment in labor and capital and high firm valuations are followed by sharp reversals (Schumpeter (1942)). During the boom, the tech sector attracts talents and especially, as we document in this paper, young talents, because sectoral choices are more easily made early in one’s career.

These early career choices are determinant for post-schooling human capital accumulation of skilled workers.¹ Therefore, the large flow of young talents in and out of the tech sector during a boom may have long-lasting consequences for productivity through its effect on skilled workers’ human capital. These consequences are beyond the direct effect of developing new technologies: They determine the different implications for long-run productivity between a smooth technological change process and the more bubbly process that characterizes technological change in the data.

One view is that workers starting in the booming tech sector are exposed to new technologies, which enables them to acquire valuable skills and enhances their long-run productivity. Even if their initial employer fails or downsizes when the tech bubble bursts, workers can redeploy these skills in other firms or other sectors. In this view, boom-bust cycles in the technology sector have positive effects on workers’ productivity, in contrast to the negative effects of boom-bust cycles in low-tech sectors such as housing (Charles, Hurst, and Notowidigdo, 2018).

The opposite view is that human capital accumulated in a booming tech sector has low value in the long run. This may happen if: workers are more likely to lose their job in the bust, thus losing firm-specific human capital in the process and facing the risk of future mismatch in the labor market; if skilled workers acquire skills that quickly depreciate due to technological acceleration; or if the boom creates a demographic imbalance in tech firms that reduces future opportunities to move up the hierarchy.

In this paper, we use administrative matched employer-employee data from France to study the long-run wage dynamics of skilled workers who start in a booming tech sector. We use as a laboratory the late 1990s Tech Bubble, during which many skilled individuals started to work in the Information and Communication Technology (ICT) sector, in order to answer three questions: First, what are the short-run and long-run wage dynamics of talents allocated to the tech sector during the boom? Second, should that long-run wage dynamics be attributed to a treatment effect of the technology boom or to selection effects or other confounding factors? Third, which economic mechanisms might explain a treatment effect of the technology boom?

We start by documenting the impact of the ICT boom for labor market allocation.

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

The share of skilled workers in the ICT sector sharply deviates from its trend between 1998 and 2001. This deviation comes almost entirely from individuals who recently entered the labor market, consistent with sectoral choices being made early in one’s career. The share of skilled labor market entrants starting in the ICT sector almost doubles during the boom, from 17.5% to 31%, before dropping back to 19% after the boom ends. The sharp delimitation in time of the boom enables us to define a “boom cohort” of skilled workers who enter the labor market during 1998–2001.

To answer the first question, we focus on the boom cohort and compare the wage dynamics of skilled workers starting in the ICT sector to that of otherwise similar individuals, but starting in a different sector. This analysis is made possible by the panel dimension of the data, which allows us to relate workers’ employment and wage dynamics to their characteristics and choices made when they entered the labor market.

Skilled workers starting in the booming ICT sector earn on average 5% higher entry wages than workers starting in other sectors. As the ICT sector contracts in the early 2000s, workers who started in ICT experience a progressive decline in their relative wage, which cancel their wage advantage by 2003. Remarkably, the relative wage of these workers keeps falling after 2003 such that, by 2015, they end up earning 6% less than workers who started outside the ICT sector. Overall, a career start in ICT is associated with 11 percentage points lower wage growth in the subsequent fifteen years or so.

We refer to this new fact as the *ICT boom-cohort discount*. It is quantitatively robust to controlling for sex, education, broad occupation at entry, regional trends, and to excluding workers starting in the financial sector from the comparison group. It is also robust to controlling for observable characteristics of workers’ initial employer that may affect workers’ earnings such as firm size, productivity or age. The ICT boom-cohort discount has a similar magnitude for workers starting in high growth firms or in subsidiaries of US companies, ruling out that it is a low quality firm or a French firm phenomenon.

Quantile regressions further show that the entire wage growth distribution of skilled workers starting in the booming ICT sector is shifted to the left. The evidence is thus inconsistent with the boom creating winners and losers among talents who go into the booming ICT sector. In particular, starting in the booming tech sector does not generate the same right-skewed distribution of earnings as becoming entrepreneur (e.g. Hamilton (2000), Kerr, Nanda, and Rhodes-Kropf (2014), Hurst and Pugsley (2015), Manso (2016)). A present value calculation shows that a career start in the booming ICT sector is associated with 4% lower cumulative earnings on average, or 18,000 euro over the first fifteen years.

Does the ICT boom-cohort discount reflect a treatment effect of the boom? We provide several pieces of evidence inconsistent with a selection effect by which workers attracted to the booming ICT sector would be of low intrinsic productivity. 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 selection hypothesis.

Second, we exploit the cohort of workers who started in the ICT sector just before the boom (between 1994 and 1996). To the extent that the boom was not anticipated in the mid-1990s, these slightly older workers were treated, but not selected by the boom. 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 the absence of negative selection for the pre-boom cohort, we find that pre-boom entrants in the ICT sector have similar or slightly higher entry wages compared to entrants in other sectors. While those workers experience a relative wage increase during the boom, their wage dynamics after the boom follows the same pattern as the boom cohort, ending up 6% below that of entrants in other sectors. The similar long-run wage dynamics of the pre-boom cohort and the boom cohort of workers starting in the ICT sector is consistent with a treatment effect of the boom, but not with a selection effect by which low quality workers would select into the booming sector.

An alternative explanation is that the ICT boom-cohort discount reflects a cycle effect, by which the ICT sector experiences an overall decline in labor productivity after the bust. We rule out this hypothesis by focusing on the cohort of workers who started in the ICT sector after the bust (between 2003 and 2005). We find that workers of the post-boom cohort who start in ICT face slightly lower entry wages than entrants in other sectors, but they catch up over time. Therefore, while both the pre-boom and boom cohorts of ICT sector entrants have poor wage growth after the boom, the pattern is reversed for ICT entrants who did not experience the boom. Overall, the evidence is consistent with a treatment effect of the boom on the long-run value of human capital of skilled workers.

Finally, we explore three potential mechanisms to explain why human capital accumulated during a technology boom has low long-term value. First, workers losing their jobs in the bust may lose firm-specific human capital or suffer scarring effects and end up on a different career path associated with long-term earnings losses (e.g. Gibbons and Katz (1991), Jacobson, LaLonde, and Sullivan (1993), von Wachter and Bender (2006) and Jarosch (2015)). Second, human capital accumulated in the ICT sector during the boom may become rapidly obsolete as a result of technological acceleration. Third, the large flow of boom-cohort workers into ICT may create a demographic imbalance that makes it difficult for workers of that cohort to move up the hierarchy in ICT firms.

We find some support for the skill obsolescence mechanism. We hypothesize that if human capital acquired by young workers during the boom depreciates rapidly because of fast technological change, then skill obsolescence should be an increasing function of the intensity of workers' job technological content. We construct three measures of job

technological intensity. First, we distinguish among skilled workers between those holding a STEM occupation (hereafter engineers) and those holding a management/business occupation (hereafter managers). Consistent with technical skills being more subject to obsolescence, we find that workers starting as engineers in the ICT sector during the boom experience low long-run wage growth, but not those starting as managers. This result also helps us rule out a selection effect, since selection does not predict that engineers only should experience an ICT boom-cohort discount, while it is a natural prediction of the skill obsolescence mechanism.

The second proxy for job technological content is firm specific and is equal to the fraction of engineers in the skilled workforce of the worker's initial employer. The third proxy is (four-digit) industry specific and is equal to the fraction of engineers in the industry in which the workers takes her first job. In both cases, the ICT boom-cohort discount is larger for workers who started in more tech-intensive firms or sectors.

We find little support for the job termination mechanism and the demographic imbalance mechanism. Regarding the former, we decompose total wage growth from entry to 2015 into a within-jobs and a between-jobs component and find that almost all of the relative wage decline takes place within jobs. We also show that controlling directly for job termination explains a negligible part of the lower wage growth. Regarding the demographic imbalance mechanism, we do not find that skilled workers from the boom cohort are less likely to be promoted.

The implications of sectoral booms on workers' choices and outcomes have been studied in other contexts. Oyer (2008) shows that MBAs are more likely to become investment bankers if the financial sector is expanding while they are in school and that earnings in investment banking are unconditionally higher than in other sectors, but he does not study how these earnings depend on market conditions when individuals start their career. Gupta and Hacamo (2018) focus on the 2000s finance boom and show that it led to a reallocation of engineers to the financial sector that made them less likely to subsequently become entrepreneurs. Charles, Hurst, and Notowidigdo (2018) analyze the 2000s housing boom and show that it reduced educational attainment because individuals dropped out of school to work in the housing sector. Choi, Lou, and Mukherjee (2017) show that the presence of salient successful firms in a sector affects college major choices and is associated with lower future wages in that sector. Our paper differs from these papers in at least two important dimensions. First, we focus on a large, well-identified technology boom. Second, we relate workers' long-term wage dynamics to their initial sectoral choice using individual panel data.

A different literature analyzes how the aggregate state of the economy when a cohort of workers enter the labor market affect their long-run outcomes. This literature finds that workers starting in a recession have persistently lower earnings (e.g. Kahn (2010), Oreopoulos, von Wachter, and Heisz (2012), Altonji, Kahn, and Speer (2016), Speer

(2016), Schwandt and von Wachter (2019)) and are less likely to reach high-end positions (e.g. Oyer (2006) for academics, Schoar and Zuo (2017) for CEOs). Our focus is different: We are interested in how sector-specific booms affect the long-run outcomes of workers allocated to the booming sector relative to same-cohort workers allocated to other sectors.

A strand of literature focuses more specifically on the late 1990s ICT boom. Beaudry, Green, and Sand (2016) argues that the overall demand for cognitive tasks declined after the tech boom. Among other differences with our paper, they do not distinguish between ICT-related tasks and other tasks and they do not distinguish between different cohorts of workers. A strand of papers argue that the high stock prices in the tech sector was a bubble (e.g. Ofek and Richardson (2003)) and facilitated investments by young tech firms (Brown, Fazzari, and Petersen (2009)) and by non-tech firms (Campello and Graham (2013)).²

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

We use several administrative data sources from France.

Exhaustive employer-employee data. 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 individuals in the private sector, with information about 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 includes unique firm and establishment identifiers that can be linked with other administrative data. We use this data set to study the composition of the labor force at the aggregate level, sector level, and firm level. The exhaustive employer-employee data does not 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 (see Abowd, Kramarz, and Margolis (1999) for a detail description of this data). We

²While not focused specifically on the 1990s boom, Dong, Hirshleifer, and Teoh (2017) show that stock market overvaluation induces firms to invest more in innovation. A different strand of literature analyzes how the composition of innovation activity varies over the business cycle (Nanda and Rhodes-Kropf (2013), Manso, Balsmeier, and Fleming (2017)).

use this panel data set to study workers' wage dynamics. An individual exits the panel 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.

Census-matched subsample. A 4/30th subsample of the panel data (individuals born in the first four days of October) can be linked with census data (*Echantillon Démographique Permanent*), which contains demographics information. We use this smaller sample to retrieve information about education.

Identifying skilled labor market entrants. We focus on the employer-employee panel over the years from 1994 to 2015. 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. Unless otherwise stated, we refer to jobs with these characteristics when we write “jobs”. 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 or work part time or over periods of less than six months.

We identify skilled workers based on occupations. The data includes a two-digit classification of job occupations (*Professions et Catégories Socioprofessionnelles*) that 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 1994–2015. Within managers and professionals, the two-digit classification distinguishes between occupations with a STEM (Science, Technology, Engineering, Mathematics) 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 A1 reports summary statistics for the sample of skilled workers over the period 1994–2015. The median skilled worker is a man (mean 0.69), is 43 year old (mean 43), and earns an annual gross salary of about 41,000 euros (mean 50,000 euros; unless otherwise stated all amounts in the paper are in constant 2000 euros).

Most of our analysis focuses on skilled individuals who enters the labor force over the period 1994–2005. We define the year of entry as the year in which the individual takes

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

her first (full-time) job, subject to the condition that she is no more than 30 year old at that time.⁵ We define skilled individuals as those holding a higher-level occupation in their first job—and not based on their current occupation because it is endogenous to human capital accumulated during the individual’s career. Appendix Table A1 reports summary statistics for skilled 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 a annual gross salary of about 38,000 euros (mean 45,000 euros).

Firm accounting data. We retrieve accounting information on firms from tax files (*FICUS-FARE*), 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 ownership data. We obtain data on firms’ ownership structure from a yearly survey of business groups (*Enquête Liaisons Financières*) run by the statistical office and crossed with information from Bureau Van Dijk. The data provides information both about direct and indirect stakes and cross-ownerships, which allows us to reconstruct group structures even in the presence of pyramids. The data includes information on the nationality of the ultimate owner, which allows us to identify subsidiaries of foreign companies.

Firm creation data. We use business creation files, which contain the list of all business creation with the date of registration. We use this data to identify startups.

3 The ICT Boom and Bust

We analyze the late 1990s boom in the Information and Communications Technology (ICT) sector using the OECD (2002) definition of ICT industries. Table A2 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. While the share of skilled workers is 15% in the overall economy, that share is

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

42% in the ICT sector. In terms of education attainment, for which we have information for a subset of individuals, the fraction of workers holding a five-year college degree is 14% over all industries, whereas that fraction is 30% in the ICT sector. The ICT sector is more specifically intensive in STEM skills: The fraction of skilled workers holding an occupation intensive in STEM skills 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. Panel A shows that the share of the ICT sector in total employment increases from just below 5% in 1996 up to 5.7% in 2001 and back to 5% by 2005. Panel B shows that the boom is more evident for skilled workers. The share of the ICT sector in skilled employment goes from 12.5% in 1996 up to 16.5% in 2001 and down to 15% in 2005. 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 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 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 from the trend during the ICT boom. By contrast, the component representing young workers exhibits a sharp upward deviation from the trend during the ICT boom. This pattern is consistent with the notion that workers can decide in which sector to start their career but inter-sector mobility after entry is not as easy. Guided by this preliminary result, we focus on labor market entrants in the rest of the paper. 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 ICT boom. It almost doubles between 1996 and 1999, from 17.5% to 31%, before dropping down to 19% in 2004.

Three results should be taken away 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 1997-8 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 and estimate the following wage equation:

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

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 job is in the ICT sector, and X_i is a vector of worker characteristics including 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 cohort and 95% confidence interval. Workers starting in the booming ICT sector earn an entry wage on average 5% higher than workers of the same cohort and with the same observable characteristics, starting outside the ICT sector. This wage difference vanishes rapidly after the boom ends in 2001. More surprisingly, the wage difference keeps falling after the bust such that by 2015, workers who started in ICT earn on average 6% less than same-cohort workers who started outside the ICT sector.

Table 1 reports the regression results. We estimate Equation (1) using for each worker, the year of entry and the years 2002, 2006, 2010, and 2015. Accordingly, the regression equation includes four coefficients β_t , measuring the wage difference between workers who started in the ICT sector and those who started in other sectors, at the time of entry, in 2002, in 2006, in 2010, and in 2015. Column 1 shows that during the boom, entrants in the ICT sector start with a wage higher by 4.6% (significant at 1%) relative to entrants in other sectors. This wage premium decreases over time and eventually becomes negative. In 2015, these workers earn on average 6.2% (significant at 1%) less than workers who started outside the ICT sector.

We include worker fixed effects in Column 2 to control for unobserved time-invariant worker heterogeneity. Worker fixed effects ensure that time variation in β_t is identified off time variation in log wage on a constant set of workers, purging potential composition effects driven by differences in propensity to exit the sample. When worker fixed effects are included, the β_t time-series is identified up to an additive constant. We use the entry year as the reference year and estimate the wage difference between entrants in the ICT sector and other entrants relative to the wage difference at entry. The pattern is similar to that without worker fixed effects: the wage difference decreases over time and reaches

−10.9% (significant at 1%) in 2015. Composition effects do not seem to be important as the relative wage discount in 2015 estimated with worker fixed effects is close to the wage discount in 2015 minus that at entry estimated without worker fixed effects.

4.2 Long-Difference Analysis

Since individuals entering into the ICT sector during the boom experience a steady wage decline after the bust, we now focus on the long difference in log wage from the entry year to 2015. We estimate the following cross-sectional regression:

$$\log(w_{i,2015}) - \log(w_{i,0}) = \beta ICT_{i,0} + \gamma X_i + \epsilon_i, \quad (2)$$

The identification of β in Equation (2) comes from the same variation in the data as the identification of β_{2015} in Equation (1) with worker fixed effects and taking the year of entry as the reference year. The coefficient on ICT_0 in Column 1 of Table 2 implies that entrants in the booming ICT sector experience 10.5 percentage points (significant at 1%) lower wage growth from entry to 2015.⁶

In the rest of the table, we rule out several basic explanations for this results. First, we control for geographical disparities in wage dynamics by adding commuting zone fixed effects in Column 2.⁷ The ICT boom-cohort discount 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 these areas during the sample period.

Second, we check whether the discount is driven by exceptionally high wage growth in a few other sectors, such as the financial sector (see Philippon and Reshef (2012) for the US and Célérier and Vallée (2017) for France). In Column 3, we exclude entrants starting in the financial sector, who represent 5% of skilled entrants during the ICT boom. The discount is slightly reduced, reflecting the high wage growth in finance during the 2000s, but it remains large and significant.

Third, we test whether the ICT boom-cohort discount is explained by observable worker characteristics. The baseline specification already controls for sex and age at entry. In Columns 4 and 5, we use the subset of the data that can be linked with census data, which provides information on the level of education. We construct two variables of educational attainment: a dummy equal to one if the individual holds at least a three-year college degree (*Licence* or equivalent) and a dummy equal to one if the individual holds at least a five-year college degree (*Master* or equivalent). 91% of skilled entrants

⁶The coefficient is not exactly equal to the one on $ICT_0 \times (t = 2015)$ in Column 2 of Table 1 because the latter depends on worker fixed effects that are estimated using the year of entry, 2002, 2006, 2010 and 2015, whereas the coefficient in Column 1 of Table 2 is estimated only using the year of entry and 2015.

⁷We define commuting zones as *départments*, which partition France into 99 areas. We obtain similar results when we use *bassins d'emplois*, which partition France into 380 areas.

hold at least a three-year college degree and 83% hold at least a five-year college degree. Column 4 shows the baseline specification on the subsample linked with census data. The discount is slightly larger than that on the main sample due to sampling noise but the difference is not statistically significant. In Column 5, we control for the level of education and this does not affect the magnitude of the discount.

Fourth, workers' earnings may be under-estimated because employer-employee data reports wages but not capital income. Capital income can be significant for entrepreneurs. It may also be relevant for employees receiving stocks or options in the firm. To account for capital income, we merge the data with firm balance sheet information and retrieve the net income of the firm. Since we do not have information on stock grants or stock options, we calculate capital income under two different assumptions. First, assuming that the CEO holds all cash flow rights, we identify the CEO using the information on occupation and we allocate the firm's net income to her.⁸ Alternatively, assuming that employees have ownership stakes in the company, we allocate the firm's net income to all skilled employees according their share in total skilled-worker wage bill. We calculate total earnings as wage plus capital income and use log of total earnings as the dependent variable. Column 6 reports the results when firm profits are allocated to the CEO only and Column 7 when they are allocated across all skilled workers. In both cases, accounting for capital income has little effect on the magnitude of the discount.

Fifth, we test whether the ICT boom-cohort discount is related to ICT employers during the boom having specific characteristics that might affect workers' long-run wage.⁹ We compare characteristics of ICT employers to that of non-ICT employers in Appendix Table A3. Panel A shows that ICT employers during the boom have on average fewer employees, are more likely to be two year old or less, and have lower value added per worker than non-ICT employers. However, these differences are not specific to the boom period. Panel B shows that ICT employers in the post-boom period (2003–2005) feature similarly different characteristics from non-ICT employers as in the boom period. In particular, differences between ICT employers and non-ICT employers during the boom are not statistically different from that after the boom, except for the probability that the employer is a startup (significant at 10%). To further check whether these differences explain the ICT boom-cohort discount, we directly control for employer characteristics in the wage growth regression. Column 1 of Table 3 shows that the wage discount is not

⁸Results are similar when we use dividends instead of net income. We prefer net income because it includes capital gains coming from undistributed profits. When the firm reports several owner-managers (one-digit occupation code 2), we split the net income equally among them.

⁹Evidence that firm characteristics have long lasting consequences on workers' earnings can be found in Garicano, Lelarge, and Van Reenen (2016) and Bloom et al. (2019) for firm size, in Ouimet and Zarutskie (2014), Burton, Dahl, and Sorenson (2017) and Babina et al. (2018) for firm age, in Abowd, Kramarz, and Margolis (1999) and Card, Heining, and Kline (2013) for firm productivity, in Tate and Yang (2015) and Cestone et al. (2017) for firms' internal labor market, in Benmelech, Bergman, and Seru (2011), Hombert and Matray (2016) and Fonseca and Doornik (2019) for credit constraints.

driven away after controlling for the initial employer’s characteristics.

Sixth, while France fully embraced the ICT revolution and produced successful ICT firms, it has not become the worldwide leader in that sector. As such, one may wonder whether the ICT boom-cohort discount is specific to workers employed by French firms or whether it also exists for US employers. The scope of this paper is limited to France, yet many large US firms have offices across the world, France included, so their workers located in France appear in our data. We use ownership data to identify subsidiaries of US companies as firms that are 100% owned by a US company. In Column 2 of Table 3, we restrict the sample to workers taking their first job in the subsidiary of a US firm. If anything, the effect is slightly larger in this subsample than in the entire sample, i.e., the ICT boom-cohort discount is not a French firm phenomenon. In a similar spirit, one may wonder whether that phenomenon originates from ICT employers with little or mild commercial success or whether it also affects individuals employed by successful firms. In Column 3, we restrict the sample to workers taking their first job in a firm that with sales growth over the next five years above 40% (the top quartile of the distribution). The ICT boom-cohort discount in this subsample of successful employers is as large as in the entire sample.

4.3 Quantiles Analysis

So far, we have shown that a career start in the booming ICT sector is associated with low average long-term wage growth. One possible interpretation is that such as career start was risky given the uncertainty regarding which firms and technologies would 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 low average wage growth may conceal a more nuanced pattern across the wage growth distribution. In particular, the low mean may be associated with 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, including the same set of control variables as in Equation (2). The main pattern that emerges is that workers starting in ICT during the boom face a wage growth discount quantitatively similar across the whole wage growth distribution, with long-run discounts ranging from 10.5% (at the 10th and 25th percentiles) to 12.1% (at the 75th percentile). If anything, the discount is larger at the top of the wage growth distribution, rejecting the hypothesis that the average discount is associated with a small probability of very positive outcomes. Thus, the boom in ICT does not appear to be a period creating winners and losers among skilled labor market entrants. Instead, it shifted the whole wage growth distribution to the left for talents who started in the booming ICT sector.

The results of the quantile regressions also suggest that the ICT boom-cohort discount is not driven by marginal workers in the booming ICT sector being of worse (unobservable) quality. Suppose it was the case, that is, suppose that the pool of workers who select into ICT during the boom consists of the set of workers who would have gone into ICT no matter what and a set of low-quality workers who select into ICT because of the boom. Such a 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 more than the top quantiles. This prediction is rejected by the quantile regressions, which show that, if anything, the top quantiles drop by more.

4.4 Cumulative Earnings

We have established that a career start in the booming ICT sector is associated with a higher wage during the boom and a lower wage after the boom. We now assess whether this leads to a higher or a lower present value of cumulative earnings from entry to 2015. For each worker, we compute cumulative earnings up to every year t post-entry by summing all the worker’s earnings from the entry year to year t discounted back to the entry year at a rate of 5% per year. We estimate Equation (2) using cumulative earnings (in log or in level) as the dependent variable.

Column 1 of Table 5 reports the result using the specification in log. It shows that skilled workers starting in ICT during the boom earn cumulative earnings from entry to 2015 that are 4.3% (significant at 1%) lower than similar workers starting in other sectors. Column 2 presents the specification in level. It estimates the discounted cumulative loss at about 18,400 euro (significant at 1%). Column 3 shows that this estimate is robust to accounting for unemployment benefits.¹⁰

5 A Treatment Effect of the Boom

In this section, we study whether the ICT boom-cohort discount reflects a treatment effect of the boom. To this aim, we investigate—and rule out—two alternative explanations.

¹⁰Since unemployment benefits (UB) are only reported starting in 2008, we assign an estimated UB when a worker 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 (see Cahuc and Prost (2015)). Since the data does 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 euro 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, or when we use a more aggressive UB length of two years if the pre-unemployment job spell lasts for at least two years.

5.1 Ruling Out a Selection Effect

The ICT boom-cohort discount may be explained by a composition effect by which the marginal worker attracted by the booming ICT sector has low intrinsic productivity. This low productivity would be masked during the boom period and become apparent when the boom ends and wages start reflecting workers' productivity more accurately. In other words, the ICT boom-cohort discount may reflect a selection effect rather than the treatment effect of starting in the ICT sector during the boom. As discussed in Section 4.3, quantile regression results are not consistent with a simple selection mechanism by which the booming ICT sector would disproportionately attract workers from the left tail of the (unobserved) productivity distribution.

To further test the selection hypothesis, we exploit the sharp delimitation of the ICT boom. Individuals entering the labor market in the period preceding the boom are unlikely to select into ICT because of the boom. Yet, this pre-boom cohort who started in ICT will experience the boom. Comparing the long-run wage dynamics of the pre-boom cohort to that of boom cohort can therefore allow us to disentangle between a treatment effect and a selection effect. Similar long-run wage dynamics would be consistent with a treatment effect. By contrast, different long-run wage dynamics would be more consistent with a selection effect.

Figure 3 shows the wage dynamics of the pre-boom cohort 1994–1996.¹¹ We estimate Equation (1) for skilled workers of that cohort and plot the time-series of β_t . The figure shows that workers starting in the ICT sector in the period preceding the ICT boom earn similar wages to that of workers starting in other sectors until the beginning of the boom. This pattern is consistent with skilled workers starting in ICT before the boom having similar intrinsic productivity to those starting in other sectors. Then, these workers experience rapid wage growth during the boom and earn at the peak of the boom a 6% wage premium on average. Crucially, when the boom ends, the pre-boom cohort experiences a similar wage dynamics to that of the boom cohort in Figure 2. The relative wage of ICT entrants declines over time. By 2015, workers who started in the ICT sector before the boom earn 6% lower wage on average relative to workers of the same cohort who started outside the ICT sector.

Regression results reported in Table 6 confirm the graphical analysis, even when worker fixed effects are included. Overall, the evidence that workers going into ICT even before the boom begins experience a very similar—qualitatively and quantitatively—long-run wage dynamics as those workers who go into ICT during the boom is at odds with a selection effect. We will show an additional piece of evidence against the selection hypothesis in Section 6.2.

¹¹We exclude 1997 from the pre-boom cohort because it might be argued that the ICT boom has already started in 1997 (see Figure 1). The results in this section are robust to including 1997 in the pre-boom cohort.

5.2 Ruling out a Cycle Effect

An alternative interpretation of our results is that wages in the ICT sector are structurally low, but it was temporarily masked by the boom and became apparent when the boom ends. However, this interpretation is not consistent with our earlier result that the wage of workers starting in the ICT sector before the boom were not lower than the starting wage in other sectors (see Figure 3). Therefore, wages are not structurally lower in ICT at least before the boom. There may still have been a break after the bust, driving down ICT sector wages.

We exploit the post-boom cohort to test for this possibility. If the ICT sector experiences an overall wage decline after the bust, then the post-boom cohort should also experience declining wages. Figure 4 shows the wage dynamics of the post-boom cohort 2003–2005.¹² Workers starting in the ICT sector after the bust have slightly lower starting wages than workers starting in other sectors. Crucially, this wage gap does not widen but, to the contrary, disappears over time.

Regression results reported in Table 7 confirm the graphical analysis. Column 1 shows that post-boom entrants starting in the ICT sector earn 2.2% (significant at 5%) lower wages than entrants in other sectors, and catch up over time such that the wage difference is small and insignificant by 2015. The specification with worker fixed effects in Column 2 yields a similar conclusion. Overall, the evidence is inconsistent with a secular decline of ICT sector wages in the wake of the ICT bust. Instead, the result that the post-boom cohort experiences an opposite wage dynamics to that of the pre-boom and boom cohort is consistent with a treatment effect affecting workers exposed to the ICT sector during the boom.

6 Economic Mechanisms

The evidence accumulated in the previous sections points towards a treatment effect of the ICT boom on the long-run value of human capital of skilled workers allocated to the booming ICT sector. We now explore three hypotheses explaining why human capital accumulated during a technology boom has low long-term value: (1) human capital is firm specific and there is a high rate of job termination when a technology boom ends; (2) human capital acquired during technology booms depreciates faster because technology changes rapidly during these periods; (3) the type of human capital acquired during technology booms is in oversupply when the booms ends. In the rest of this section, we provide evidence consistent with (2) but not with (1) or (3).

¹²We exclude 2002 from the post-boom period in order to leave a gap year between the boom period and post-boom period. The results in this section are robust to including 2002 in the post-boom period.

6.1 Job Termination

Within-jobs/between-jobs wage growth decomposition. As an elementary test of the job loss channel, we focus on the boom cohort and decompose workers' wage growth from entry to 2015 into a within-jobs and a between-jobs components. If lower wage growth is explained (economically) by workers forced to change jobs and in the process losing human capital, then it should be explained (statistically) by the between-jobs component. Indexing by $t = 0, \dots, T$ the years in which we observe worker i and denoting by $F_{i,t}$ her employer in year t , we construct within-jobs wage growth as

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

and between-jobs wage growth as

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

We estimate Equation (2) using these two components of wage growth as dependent variables.

Table 8 shows that the wage growth discount comes almost entirely from the within-jobs component. Of the total 10.5% lower wage growth experienced by entrants in the booming ICT sector, 8.8 percentage points (significant at 1%) come from lower wage growth within job spells and only 1.7 percentage points (insignificant) come from lower wage growth during job transitions. This result does not reflect the fact that wage growth happens almost only within job spells unconditionally: for skilled entrants (in any sector) during the boom period, within-jobs and between-jobs wage growth explain respectively 39% and 18% of the variation in total wage growth.¹³

The decomposition into a within/between-jobs may still underestimate the effect of job termination if job termination reduces the probability of promotion or increases the risk of mismatch in the new job, thereby weighing on future (within-jobs) wage growth. To address this possibility, we now test directly whether (forced) job termination explains the ICT boom-cohort discount.

Job termination. We construct four variables to measure job termination. The first two do not distinguish between forced and voluntary job termination: (1) a dummy variable equal to one if the worker changes employer within the first four years after entry; and (2) a dummy variable equal to one if the worker has changed employer by 2015. The next two are dummy variables equal to one if the worker experiences a forced

¹³The R^2 do not sum to one because within-jobs wage growth and between-jobs wage growth are negatively correlated in the cross-section of workers.

job termination within the first four years after entry, where forced termination is defined as (3) a transition to a lower-paid job with another employer or as (4) a transition to another employer when the initial employer has negative employment growth in the year of the transition. The unconditional probability of job termination is, for each of the four proxies, 59%, 86%, 17%, and 20%, respectively.

Table 9 shows how the probability of job termination depends on the sector of entry for the pre-boom, boom, and post-boom cohorts. We regress each of the job termination dummy on ICT_0 interacted with dummy variables for the pre-boom cohort (entry year during 1994–1996), the boom cohort (1998–2001), and the post-boom cohort (2003–2005), and the same set of controls as before, all interacted with the cohort dummies. When we consider all types of job termination in Columns 1 and 2, skilled workers starting in the ICT sector during the boom are more likely to experience job termination than those of the pre-boom cohort but not more than those of the post-boom cohort.

When we focus on forced job termination in Columns 3 and 4, a clear pattern emerges. Workers starting in ICT during the boom are more likely to experience forced termination than workers starting in ICT before or after the boom. This result holds for both proxies of forced termination. For instance, Column 3 shows that ICT entrants during the boom are 4.6 percentage points (significant at 1%) more likely to experience a transition to a lower-paid job within the first four years of their career than entrants in other sectors. By contrast, there is no significant difference for the pre-boom cohort and the post-boom cohort.¹⁴

We go on testing whether the higher probability of job loss explains the ICT boom-cohort discount. We re-estimate Equation (2) for the boom cohort controlling directly for each of the four proxies of job termination. The odd-numbered columns of Table 10 show that job termination explains a negligible part of the discount. Compared to the baseline discount of 10.5% (Column 1 of Table 2), job termination explains at most 0.7 percentage points of this discount (using the first proxy of forced job termination in Column 5).¹⁵

Job termination during a sectoral bust might impact disproportionately the wage. The specification in the even-numbered columns of Table 10 includes an interaction term between ICT_0 and job termination 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 (non-interacted) ICT_0 can be interpreted as the wage growth difference between workers starting in the ICT sector and experiencing no job termination, and entrants in other sectors experiencing no job termination. With all four proxies of job termination, we find a wage growth discount of the same magnitude (in the range 8.3% to 11.4%) as in the baseline specification (10.5%). A particularly telling result is the

¹⁴The difference in coefficient between the boom cohort and pre-boom cohort is significant at 1% and the one between the boom cohort and post-boom cohort is significant at 5%.

¹⁵A similar conclusion obtains when we include all four proxies of job termination in the same regression. In this specification, job termination explains 0.5 percentage points of the discount.

one reported in Column 4, showing that workers starting in the ICT sector during the boom and still working with their initial employer in 2015 experience a 8.3 percentage points lower wage growth than entrants in other sectors and also working with their initial employer in 2015.

Overall, the evidence is inconsistent with the ICT boom-cohort discount being explained by job losses. Even workers who do not switch employers face as poor a long-term wage growth as those losing their jobs. The evidence suggests that these workers experience a long-run decline in productivity that goes beyond firm-specific effects and that materializes regardless of their career path.

6.2 Skill Obsolescence

Skilled workers starting in ICT during the boom may accumulate human capital early in their career that rapidly becomes obsolete because technology evolves fast during technology booms. If the ICT boom-cohort discount is explained by this mechanism, we expect it to be larger for workers holding a job with a higher technological content or working in firms more intensive in technology, because 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 is constructed using the occupation held by the worker at entry. The two-digit occupation classification in the data distinguishes between occupations with a science/engineering skill content (hereafter "engineers") and those with a management/business content (hereafter "managers"). We define *Engineer* as the dummy variable equal to one if the worker holds an engineering occupation in her first job. The second proxy aims at capturing the technological intensity of firms in which skilled workers start their career. We define *TechFirm* as the fraction of engineers in the worker's initial employer. The third proxy aims at capturing the technological intensity of specific (four-digit) sectors of the broad ICT sector in which workers start their career. We define *TechSector* as the fraction of engineers in the four-digit sector in which the worker 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 between ICT_0 and the engineer dummy as an explanatory variable.¹⁶ The coefficient on the interaction term shows that engineers who started in ICT have 9.9 percentage points (significant at 5%) lower wage growth than managers who started in ICT and relative to the same difference in other sectors. By contrast, the coefficient on the non-interacted ICT dummy is small and insignificant, showing that managers starting in ICT do not have lower wage growth than managers starting in other sectors. Thus, consistent with the skill obsolescence

¹⁶The non-interacted variable *Engineer* is not included because the baseline specification already has fixed effects for the initial occupation.

hypothesis, the long-run wage discount is concentrated on engineers.

Next, we turn to the impact of the initial employer’s technological intensity. In Column 2, we include the interaction of ICT_0 with $TechFirm$ as an explanatory variable. The coefficient on the interaction term is negative and significant at the 1% level. Thus, the discount is stronger for workers who started in more-tech firms. One concern could be that the result is driven by a more general pattern by which skilled workers starting in more-tech firms even outside the ICT sector would experience lower wage growth. We test this hypothesis in Column 3 by adding the dummy $(1 - ICT_0)$ interacted with $TechFirm$. Two results appear. First, the impact of the firm’s technological intensity for workers starting in ICT is barely affected by the inclusion of that variable. Second, the firm’s technological intensity has no significant impact for workers starting outside ICT. Thus, patterns of wage dynamics are 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.

A similar pattern emerges when we use the proxy for the sector’s technological intensity. The top three ICT industries 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”. Column 4 shows that the ICT boom-cohort discount is stronger for workers 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.

On a different note, the evidence that engineers but not managers experience the ICT boom-cohort discount goes against a selection mechanism by which individuals with low unobserved ability would select into the booming ICT sector, as discussed in Section 5.1. Indeed, there is no clear reason why engineers’ sectoral choice would be more responsive to market conditions than managers’. If anything, engineers might have more specific skills and thus be more constrained in their sectoral choice. By contrast, the skill obsolescence mechanism naturally generates the prediction that engineers should be more affected by the ICT boom than managers.

6.3 Demographic Imbalance

Large inflows of skilled workers into ICT during the boom may have led to an oversupply of workers from the boom cohort. If, as argued for instance by Welch (1979) or Jeong, Kim, and Manovskii (2015), workers from different cohorts are imperfect substitutes, then the demographic imbalance created by the boom may explain why wage growth is low for the ICT boom cohort but not for the post-boom cohort. One mechanism by which workers of the boom cohort and workers of the post-boom cohort could be complements

rather than substitutes, is that experienced skilled workers become managers of junior skilled workers. A distorted age pyramid in the ICT sector may create a bottleneck that makes it less likely for the boom-cohort workers to be promoted.

To test this hypothesis, we focus on workers starting as engineers and analyze whether these workers are less likely to be promoted manager if they start in the ICT sector during the boom. The sample for this test is skilled workers from the boom cohort (1998–2001) and the post-boom cohort (2003–2005) who start as engineers based on the two-digit occupation classification described in Section 6.2. We construct a dummy variable *Promotion* equal to one if the worker has become a manager in her starting industry in 2015.¹⁷ The unconditional mean of the promotion dummy is 0.32. To validate the proxy for promotion, we regress wage growth from entry to 2015 on the promotion dummy and the same set of controls as before and four-digit industry fixed effects. Column 1 of Table 12 shows that engineers who follow a career path leading up to a management position in their starting industry experience a 12% (significant at 1%) higher wage growth than engineers who follow a different career path.¹⁸ In Column 2, we interact the promotion dummy with ICT_0 and find that the interaction term is small and insignificant. Thus, the proxy for promotion is similarly valid for engineers starting in or outside the ICT sector.

We go on testing whether engineers starting in ICT during the boom have a lower probability of being promoted. We follow a difference-in-difference approach and compare the probability of promotion for engineers who started in the ICT sector relative to engineers who started in other sectors (first difference) for the boom cohort relative to the post-boom cohort (second difference). We regress the promotion dummy on ICT_0 interacted with a boom cohort dummy. In Column 3, the coefficient on the interaction term is small and statistically insignificant.¹⁹ Therefore, the large flow of engineers to ICT during the boom does not seem to have reduced these workers' future opportunities of promotion to management positions.

¹⁷We use a broad industry classification (with 10 different industries) to determine whether the worker has become a manager in the same industry in which she started her career. We obtain similar results if we use the two-digit industry classification (84 industries) or the four-digit industry classification (476 industries).

¹⁸Only a small part of this effect comes from the fact that engineers who still work in their starting industry in 2015 have higher wage growth even if they do not become managers. When we add a dummy variable equal to one if the worker still works in her starting industry in 2015, the coefficient on this dummy variable is 0.021 (significant at 10%) and the coefficient on the promotion dummy is 0.10 (significant at 1%).

¹⁹Columns 2 and 3 do not include non-interacted ICT_0 or boom cohort dummy as explanatory variables because the specifications already include industry fixed effects and entry year fixed effects.

7 Conclusion

In this paper, we uncover a new fact that we call: the *ICT boom-cohort discount*. Young talents who started in the ICT sector during the late 1990s Tech Bubble enjoyed 5% higher entry wages, but end up in the long run with 6% lower wages, relative to similar skilled workers who started in a different sector. This finding is not consistent with the notion that boom-time tech firms enhance their workers' human capital and long-term productivity. One question left unanswered by our paper is that of the impact of these firms on aggregate productivity, as they may create positive spillovers on workers in other sectors or on future innovation. We leave this question for future research.

References

- Abowd, John M, Francis Kramarz, and David N Margolis. 1999. “High Wage Workers and High Wage Firms.” Econometrica 67 (2):251–333.
- 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.
- Babina, Tania, Ma Wenting, Ouimet Paige, and Zarutskie Rebecca. 2018. “Entrepreneurial Wages.” Working Paper .
- 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.
- Benmelech, Efraim, Nittai K Bergman, and Amit Seru. 2011. “Financing Labor.”
- Bloom, Nicholas, Guvenen Fatih, Lucking Brian, Smith Ben, Song Jae, and von Wachter Til. 2019. “The Large Firm Pay Premium Redux.” Working Paper .
- 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.
- Burton, M Diane, Michael S Dahl, and Olav Sorenson. 2017. “Do Start-ups Pay Less?” Industrial and Labor Relations 71 (5):1179–1200. URL <https://doi.org/10.1177/0019793917747240>.
- Cahuc, Pierre and Corinne Prost. 2015. “Améliorer l’assurance chômage pour limiter l’instabilité de l’emploi.” Notes du Conseil d’Analyse Economique 24:1–12.
- 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.
- Card, David, Jörg Heining, and Patrick Kline. 2013. “Workplace Heterogeneity and the Rise of West German Wage Inequality.” Quarterly Journal of Economics 128 (3):967–1015. URL <https://www.jstor.org/stable/26372517>.
- Célérier, Claire and Boris Vallée. 2017. “Returns to talent and the finance wage premium.” Working Paper .
- Cestone, Giacinta, Chiara Fumagalli, Francis Kramarz, and Giovanni Pica. 2017. “Insurance Between Firms : The Role of Internal Labor Markets.” Working Paper .

- Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo. 2018. “Housing booms and busts, labor market opportunities, and college attendance.” American Economic Review 108 (10):2947–94.
- Choi, Darwin, Dong Lou, and Abhiroop Mukherjee. 2017. “The effect of superstar firms on college major choice.” Working Paper .
- Dong, Ming, David Hirshleifer, and Siew Hong Teoh. 2017. “Stock Market Overvaluation, Moon Shots, and Corporate Innovation.” NBER Working Paper .
- Fonseca, Julia and Bernardus Van Doornik. 2019. “Financial Development , Labor Markets , and Aggregate Productivity: Evidence from Brazil.” Working Paper .
- Garicano, Luis, Claire Lelarge, and John Van Reenen. 2016. “Firm Size Distortions and the Productivity Distribution: Evidence from France.” American Economic Review 106 (11):3439–3479. URL <http://www.jstor.org/stable/24911312>.
- Gibbons, Robert and Lawrence F Katz. 1991. “Layoffs and Lemons.” Journal of Labor Economics 9 (4):351–380. URL <http://www.jstor.org/stable/2535075>.
- Gibbons, Robert and Michael Waldman. 2006. “Enriching a theory of wage and promotion dynamics inside firms.” Journal of Labor Economics 24 (1):59–107.
- Gupta, Nandini and Isaac Hacamo. 2018. “Superstar (and Entrepreneurial) Engineers in Finance.” Working Paper (March).
- Hamilton, Barton H. 2000. “Does entrepreneurship pay? An empirical analysis of the returns to self-employment.” Journal of Political economy 108 (3):604–631.
- Hombert, Johan and Adrien Matray. 2016. “The Real Effects of Lending Relationships on Innovative Firms and Inventor Mobility.” Review of Financial Studies 30 (7):2413–2445. URL <https://dx.doi.org/10.1093/rfs/hhw069>.
- 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 .
- Jeong, Hyeok, Yong Kim, and Iourii Manovskii. 2015. “The price of experience.” American Economic Review 105 (2):784–815.

- 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. 2016. “Experimentation and the Returns to Entrepreneurship.” Review of Financial Studies 29 (9):2319–2340. URL <http://www.jstor.org/stable/2999586>.
- 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.” Journal of Finance 58 (3):1113–1137. URL <https://doi.org/10.1111/1540-6261.00560>.
- 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.
- Quimet, 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.
- Oyer, Paul. 2006. “Initial Labor Market Conditions and Long-Term Outcomes for Economists.” Journal of Economic Perspectives 20 (3):143–160.
- . 2008. “The Making Stock Market of an Investment Banker: Shocks, Career Choice, and Lifetime Income.” Journal of Finance 63 (6):2601–2628.
- 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.
- Schoar, Antoinette and Luo Zuo. 2017. “Shaped by Booms and Busts: How the Economy Impacts CEO Careers and Management Styles.” Review of Financial Studies 30 (5):1425–1456. URL <http://dx.doi.org/10.1093/rfs/hhw111>.
- Schumpeter, Joseph A. 1942. Capitalism, socialism and democracy. Harper & Brothers.

Schwandt, Hannes and Till von Wachter. 2019. “Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets.” Journal of Labor Economics 37 (S1):S161–S198. URL <https://doi.org/10.1086/701046>.

Speer, Jamin. 2016. “Wages, Hours, and the School-to-Work Transition: The Consequences of Leaving School in a Recession for Less-Educated Men.” B.E. Journal of Economic Analysis & Policy 16:97. URL <https://www.degruyter.com/view/j/bejeap.2016.16.issue-1/bejeap-2015-0054/bejeap-2015-0054>.

Tate, Geoffrey and Liu Yang. 2015. “The Bright Side of Corporate Diversification: Evidence from Internal Labor Markets.” Review of Financial Studies 28 (8):2203–2249.


von Wachter, Till and Stefan Bender. 2006. “In the Right Place at the Wrong Time: The Role of Firms and Luck in Young Workers’ Careers.” American Economic Review 96 (5):1679–1705. URL <http://www.jstor.org/stable/30034990>.

Welch, Finis. 1979. “Effects of cohort size on earnings: The baby boom babies’ financial bust.” Journal of Political Economy 87 (5, Part 2):S65–S97.


Figure 1: Employment Share of the ICT Sector

Panel A shows the share of the ICT sector in total employment. Panel B shows the share of the ICT sector in skilled employment. 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.


Panel A: All workers



Panel B: Skilled workers



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



Panel D: Skilled entrants




Figure 2: Wage Dynamics of the ICT Boom Cohort

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. Dashed lines represent the 95% confidence interval. The regression is estimated over the cohort of skilled workers whose first full-time job was in 1998–2001.



Figure 3: Wage Dynamics of the Pre-Boom Cohort

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. Dashed lines represent the 95% confidence interval. The regression is estimated over the cohort of skilled workers whose first full-time job was in 1994–1996.




Figure 4: Wage Dynamics of the Post-Boom Cohort

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. Dashed lines represent the 95% confidence interval. The regression is estimated over the cohort of skilled workers whose first full-time job was in 2003–2005.



Table 1: 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 equal to one if worker i started in the ICT sector. $(t=Y)$ is a dummy equal to one if year t is Y = entry year, 2002, 2006, 2010, or 2015. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. 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)
$(ICT_0) \times (t=0)$.046*** (.007)	
$(ICT_0) \times (t=2002)$.030*** (.007)	-.004 (.007)
$(ICT_0) \times (t=2006)$	-.025** (.010)	-.070*** (.001)
$(ICT_0) \times (t=2010)$	-.051*** (.012)	-.095*** (.011)
$(ICT_0) \times (t=2015)$	-.062*** (.015)	-.109*** (.014)
Worker controls	✓	✓
Worker FE	–	✓
Observations	31,670	30,423

Table 2: 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 growth of worker i from entry year to 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. From Column 2 on, Commuting Zone fixed effects are included. In Column 3, entrants who started in the finance sector are excluded. In Column 4, the sample is restricted to workers that can be linked with census data. In Column 5, we add two dummy variables for the worker holding a three-year college degree and for the worker holding a five-year college degree. In Column 6, the firm’s net income is added to the worker’s wage if the worker is the CEO of the firm. In Column 7, a fraction of the firm’s net income equal to the worker’s share in total wage bill is added to the worker’s wage. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2015 – log wage entry						
	(1)	(2)	(3)	(4)	(5)	Incl. firm profit for CEO	all emp.
	(6)	(7)					
ICT_0	-.105*** (.015)	-.113*** (.016)	-.104*** (.016)	-.154*** (.044)	-.152*** (.043)	-.133*** (.016)	-.130*** (.016)
Worker controls	✓	✓	✓	✓	✓	✓	✓
Commuting Zone FE	–	✓	✓	✓	✓	✓	✓
Education	–	–	–	–	✓	–	–
Observations	4,972	4,972	4,599	537	537	4,948	4,968
Sample	All	All	Excl. finance	Census	Census	All	All

Table 3: Wage Growth and Firm Characteristics

The table presents OLS estimations of Equation (2) for skilled entrants of the boom cohort 1998–2001. The dependent variable is wage growth growth of worker i from entry year to 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. $\text{Log}(\text{Employees})$, $\text{Value added/Worker}$, and Startup are variables defined for the initial employer of worker i and equal to the log number of employees, value added per worker, and a dummy equal to one if the firm is two year old or less, respectively. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. In Column 2, we restrict the sample to workers whose initial employer is the subsidiary of a US company. In Column 3, we restrict the sample to workers whose initial employer has sales growth in the subsequent five years above 40%. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2015 – log wage entry		
	(1)	(2)	(3)
ICT_0	-0.11*** (.015)	-0.15*** (.045)	-0.095*** (.032)
$\text{Log}(\text{Employees})$	0.0026 (.0032)		
$\text{Value added/Worker}$	0.00085*** (.00015)		
Startup	0.042 (.026)		
Worker controls	✓	✓	✓
Observations	4,282	530	879
Sample	All	US firms	High growth firms

Table 4: 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 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, 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	-.105*** (.026)	-.105*** (.019)	-.107*** (.010)	-.121*** (.018)	-.110*** (.029)
Worker Controls	✓	✓	✓	✓	✓
Observations	4,972	4,972	4,972	4,972	4,972

Table 5: Cumulative Earnings

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 discounted cumulative earnings of worker i from entry year to year t , in log in Column 1 and in level in Column 2. In Column 3, earnings include unemployment benefits assuming a 60% replacement rate for one year. ICT_0 is a dummy equal to one if worker i started in the ICT sector. $(t=Y)$ is a dummy equal to one if year t is $Y =$ entry year, 2002, 2006, 2010, or 2015. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Cumulative Earnings		
	Log (1)	Level (in Euro) (2)	Level (in Euro) incl. UB (3)
$(ICT_0) \times (t=0)$.038*** (.008)	810 *** (222)	810*** (222)
$(ICT_0) \times (t=2002)$.023*** (.011)	1748 (949)	2060** (923)
$(ICT_0) \times (t=2006)$	-.003 (.015)	-948 (2184)	-1260 (2155)
$(ICT_0) \times (t=2010)$	-.024 (.018)	-8393** (3702)	-9016** (3664)
$(ICT_0) \times (t=2015)$	-.043*** (.021)	-18381*** (5968)	-19387*** (5946)
Worker controls	✓	✓	✓
Observations	45,695	45,695	45,695

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 equal to one if worker i started in the ICT sector. $(t=Y)$ is a dummy equal to one if year t is $Y =$ entry year, 1997, 2000, 2002, 2006, 2010, or 2015. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. 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)
$(ICT_0) \times (t=0)$.028*** (.0098)	
$(ICT_0) \times (t=1997)$.018* (.011)	.0017 (.0096)
$(ICT_0) \times (t=2000)$.067*** (.015)	.056*** (.014)
$(ICT_0) \times (t=2002)$.04** (.016)	.028** (.014)
$(ICT_0) \times (t=2006)$	-.024 (.019)	-.041** (.018)
$(ICT_0) \times (t=2010)$	-.052** (.021)	-.063*** (.019)
$(ICT_0) \times (t=2015)$	-.062** (.024)	-.086*** (.022)
Worker controls	✓	✓
Worker FE	–	✓
Observations	24,546	23,403

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 post-boom cohort 2003–2005. The dependent variable is log wage of worker i in year t . ICT_0 is a dummy equal to one if worker i started in the ICT sector. $(t=Y)$ is a dummy equal to one if year t is $Y =$ entry year, 2006, 2010, or 2015. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. 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)
$(ICT_0) \times (t=0)$	-.022** (.0096)	
$(ICT_0) \times (t=2006)$	-.02* (.011)	.0091 (.0092)
$(ICT_0) \times (t=2010)$	-.002 (.014)	.026** (.012)
$(ICT_0) \times (t=2015)$.0036 (.019)	.027 (.017)
Worker controls	✓	✓
Worker FE	–	✓
Observations	15,424	14,815

Table 8: Within-Jobs/Between-Jobs Wage Growth Decomposition

The table presents the decomposition of workers' wage growth from entry to 2015 into a within-jobs component and a between-jobs component as defined in Equations (3) and (4), for skilled entrants of the boom cohort 1998–2001. ICT_0 is a dummy equal to one if worker i started in the ICT sector. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2015 – log wage entry	
	Within-jobs (1)	Between-jobs (2)
ICT_0	-.088*** (.015)	-.017 (.013)
Worker controls	✓	✓
Observations	4,972	4,972

Table 9: Job Termination

The table presents OLS regressions for skilled entrants of the pre-boom cohort 1996-1998, boom cohort 1998–2001, and post-boom cohort 2003–2005. The dependent variable is a dummy equal to one if worker i experiences job termination. In Column 1, job termination equals one if the worker switches job within the first four years after entry. In Column 2, job termination equals one if the worker has a different employer in 2015 than at entry. In columns 3, job termination equals if the worker switches job during the first four years after entry and this switch is associated with a wage drop. In columns 4, job termination equals if the worker switches job during the first four years after entry and the initial employer has negative employment growth in the year of the switch. ICT_0 is a dummy equal to one if worker i started in the ICT sector. Pre-boom cohort, Boom cohort, and Post-boom cohort are dummy variables equal to one if the worker enters the labor market over 1994–1996, 1998–2001, and 2003–2005 respectively. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	=1 if job terminated			
	forced or voluntary		forced	
	Within four years (1)	Diff. employer in 2015 (2)	Within 4y & $\Delta wage < 0$ (3)	Within 4y & $\Delta emp < 0$ (4)
(Pre-boom cohort) \times (ICT_0)	.052** (.024)	-.0081 (.016)	-.008 (.019)	-.024 (.021)
(Boom cohort) \times (ICT_0)	.076*** (.016)	.058*** (.010)	.046*** (.013)	.028** (.014)
(Post-boom cohort) \times (ICT_0)	.084*** (.024)	.058*** (.018)	-.001 (.019)	-.003 (.021)
Worker controls	✓	✓	✓	✓
Observations	10,464	10,464	10,464	10,464

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 growth of worker i from entry year to 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. In odd-numbered columns, we include each of the four proxies for job termination used in Table 9 as an explanatory variable. In even-numbered columns, we also include the interaction between ICT_0 and the proxy for job termination. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

<i>Proxy for job termination:</i>	Log wage 2015 – log wage entry							
	Within four years		Diff. employer in 2015		Within 4y & Δ wage<0		Within 4y & Δ emp<0	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ICT_0	-.102*** (.015)	-.114*** (.022)	-.102*** (.015)	-.083** (.042)	-.098*** (.015)	-.099*** (.016)	-.104*** (.015)	-.114*** (.016)
Job termination	-.035*** (.013)	-.04** (.016)	-.054*** (.018)	-.049** (.021)	-.14*** (.017)	-.15*** (.022)	-.028* (.017)	-.043** (.022)
$ICT_0 \times$ Job termination		.019 (.027)		-.021 (.043)		.005 (.035)		.041 (.034)
Worker controls	✓	✓	✓	✓	✓	✓	✓	✓
Observations	4,972	4,972	4,972	4,972	4,972	4,972	4,972	4,972

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 growth of worker i from entry year to 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. $Engineer$ is a dummy equal to one if worker i has a science/engineer (as opposed to management/business) occupation in her first job. $TechFirm$ is the fraction of engineers in worker i 's initial employer. $TechSector$ is the fraction of engineers in worker i 's initial four-digit industry. Worker controls include sex, age and age squared at entry, entry year, and two-digit occupation at entry. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2015 – log wage entry				
	(1)	(2)	(3)	(4)	(5)
ICT_0	-.027 (.039)	-.044 (.032)	-.05 (.034)	-.035 (.04)	-.043 (.042)
$ICT_0 \times Engineer$	-.099** (.042)				
$ICT_0 \times TechFirm$		-.11*** (.043)	-.12*** (.043)		
$(1 - ICT_0) \times TechFirm$			-.032 (.036)		
$ICT_0 \times TechSector$				-.17** (.083)	-.17** (.083)
$(1 - ICT_0) \times TechSector$					-.082 (.086)
Worker controls	✓	✓	✓	✓	✓
Observations	4,972	4,897	4,897	4,970	4,970

Table 12: Promotion Up the Hierarchy

The table presents OLS regressions for skilled entrants of the boom cohort 1998–2001 and post-boom cohort 2003–2005. In Columns 1 and 2, the dependent variable is wage growth growth of worker i from entry year to 2015. Promotion is a dummy equal to one if worker i has become a manager in her initial industry in 2015. ICT_0 is a dummy equal to one if worker i started in the ICT sector. In Column 3, the dependent variable is the promotion dummy. Boom cohort is a dummy equal one if the worker enters the labor market over 1998–2001. Worker controls include sex, age and age squared at entry, entry year, two-digit occupation at entry, and four-digit industry fixed effects. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Log wage 2015 – log wage entry		=1 if Promotion
	(1)	(2)	(3)
Promotion	.12*** (.019)	.14*** (.023)	
Promotion \times ICT_0		-0.29 (.041)	
$ICT_0 \times$ Boom Cohort			-.027 (.029)
Worker controls	✓	✓	✓
Industry FE	✓	✓	✓
Observations	4,228	4,228	4,228

Appendix

Table A1: Summary Statistics

Panel A shows summary statistics at the worker-year level for the period 1994–2015 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 over 1994–2005.

	N	Mean	P25	P50	P75
<i>Panel A: All skilled workers</i>					
Annual wage	1,980,099	50,406	32,137	41,414	56,468
Male	1,980,099	0.69	0	1	1
Age	1,980,099	43	35	43	51
<i>Panel B: Skilled workers entering the labor force over 1994–2005</i>					
Annual wage	244,139	44,770	29,769	38,331	50,962
Male	244,139	0.68	0	1	1
Age at entry	244,139	26	25	26	27

Table A2: ICT Industries

List of ICT industries from OECD (2002). The third (fourth) column reports the 1994–2008 average share in total employment (in skilled employment) of each ICT industry.

ICT industries	ISIC rev 3.1 codes	Share of total employment (%)	Share of skilled employment (%)
ICT: Services		1.8	7.6
IT consultancy	<i>7210</i>	0.7	3.4
Software	<i>7220</i>	0.7	3.1
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		1.2	2.1
Telecommunications	<i>6420</i>	1.2	2.1
ICT: Manufacturing		1.6	3.7
Electronic/communication equipment	<i>3210, 3220, 3230</i>	0.8	1.7
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		5.1	14.6

Table A3: Employers' Characteristics

The table reports summary statistics on the characteristics of the employers of skilled labor market entrants in the ICT sector (column 1) and in other sectors (column 2) over 1998–2001 (Panel A) and over 2003–2005 (Panel B). Column 3 reports the difference between column 1 and column 2. Employees is the number of full-time equivalent employees. Value added/Worker is value added in thousand euro per worker. Startup is a dummy equal to one if the firm is two year old or younger. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	ICT firms (1)	Non-ICT firms (2)	(1) minus (2) (3)
<i>Panel A: Boom cohort</i>			
Log(Employees)	5	5.2	-.24**
Value added/Worker	61	67	-5.5**
Startup	.15	.074	.074***
<i>Panel B: Post-boom cohort</i>			
Employees	4.8	5.1	-.24*
Value added/Worker	66	70	-4.8*
Startup	.089	.05	.039**