Not All That Glitters is Gold: Firm Hiring in the Market for Knowledge Workers*

Draft only - Do not distribute

Justine Boudou
Harvard Business School
jboudou@hbs.edu

Abstract

Firms increasingly compete for scientific talent to drive innovation. While individuals' productivity in an industrial setting is unknown at the time of hire, firms can observe measures of academic productivity. However, because Academia and Industry operate under different institutional logics, the scientists most productive in Academia may not be those most productive in Industry. Using confidential data on 40 cohorts of U.S. PhDs linked to wage and publication records, I estimate individual-level counterfactual productivity in both Academia and Industry. I do so using Marginal Treatment Effects and an instrument based on variation in firms' demand for PhDs upon graduation. I find that individuals' academic and industrial productivity are only moderately aligned (correlation=0.28), and that this alignment has declined over time. In fields with weaker alignment, firms are more likely to hire candidates who are subsequently laid off – despite these individuals displaying strong measures of academic productivity at graduation. I then estimate how much firms would need to pay to hire scientists and find substantial heterogeneity: 40% would require nearly five times the academic wage to work in Industry, while 10% would accept lower or equal pay. Importantly, individuals with higher industrial productivity require lower premiums. Taken together, the findings suggest that reliance on academic indicators can lead to costly hiring mismatches for firms. Overall, this study underscores the challenges firms face when hiring talent across institutional boundaries.

Keywords: *Human Capital*, *Innovation*, *Hiring*, *Knowledge Workers*, *Academia*, *Industry*

*The use of confidential NSF data does not imply NSF endorsement of the research methods or conclusions contained in this report.

1 Introduction

Knowledge workers play a central role in firm innovation (Kaiser et al., 2018). In knowledge-intensive industries such as artificial intelligence or biotechnology, firms increasingly compete for these highly educated individuals who possess scientific and technical expertise (Hiraiwa et al., 2025). This is reflected in significant changes in the allocation of highly trained talent: in the United States, for example, the share of PhDs in Computer Science entering Industry rose from 33% in 2003 to 70% in 2023. This shift extends beyond early-career hires, as firms also recruit senior faculty and, in some cases, entire research teams from academic institutions (Yue, 2024). Despite the substantial costs at stake, hiring scientific talent is challenging because firms do not know how productive individuals will be in an industrial setting at the time of hire (Campbell, Coff and Kryscynski, 2012). Scientists are rooted in academic environments through their training or professional experience, and firms must infer their potential from measures of academic productivity, such as publications or institutional prestige (Merton, 1973).

Academia and Industry, however, operate under different institutional logics, characterized by distinct objectives, performance metrics, and incentive structures (Gittelman and Kogut, 2003; Perkmann et al., 2013; Stern, 2004; Stuart and Ding, 2006). Academia rewards intellectual autonomy and novelty, while Industry emphasizes execution under constraints and a focus on product-oriented outcomes (Bikard, 2018). Hence, while some skills such as technical expertise or creativity are valued in both sectors, others are emphasized in one context but not the other. For example, risk-taking and team execution tend to be critical in Industry but less emphasized in an academic setting, whereas theoretical depth and independent inquiry are more important in Academia (Agarwal and Ohyama, 2013; Sauermann and Stephan, 2013). As a result, the scientists who would be the most productive in Academia might not be those who would be the most productive in Industry.

This misalignment may create challenges in selecting suitable candidates from the applicant pool and attracting high-potential individuals into it. First, if academic and industrial productivity are poorly aligned, relying on academic productivity to assess individuals' counterfactual productivity in Industry may lead firms to overlook candidates with strong potential in Industry, or to hire those whose skills are poorly matched to industrial needs (Bidwell, 2011). Second, firms must design offers that successfully appeal to individuals with strong potential in Industry, without overpaying for those who have strong academic productivity that does not translate into productivity in Industry. This also relies on understanding how productivity aligns across contexts, so that firms can identify whom to target and how much to offer (Roach and Sauermann, 2024; Stern, 2004).

In this paper, I ask: *To what extent does academic productivity reflect a scientist's potential in an industrial setting? And how does this impact firms' ability to hire high-performing scientific talent?*

Conceptually, I characterize each individual by two productivity parameters: their productivity in an academic setting, and their productivity in an industrial setting. Because different sectors may reward different skills, individuals may be highly productive in one sector but not in the other. This framework decouples individuals' potential productivity in each sector from their actual sorting outcomes, allowing me to characterize the pool of talent available to firms at the time of hire. When academic and industrial productivity are strongly correlated, academic productivity serves as a reliable proxy for industrial productivity, enabling firms to make informed decisions about whom to select and target. But when the correlation is weak or negative, academic productivity becomes noisy or misleading. This increases the risk that firms may select candidates with low industrial productivity, or design incentives that attract individuals with strong academic credentials but limited potential in Industry. Theoretically, the correlation between academic and industrial productivity is ambiguous. If Industry and Academia primarily value the same skills, the correlation will be positive. If sector-specific skills dominate, the correlation will be weaker or even negative. While existing work emphasizes the divergent institutional logics of Academia and Industry (Stephan, 2012), others argue that sectoral differences may be overstated (Sauermann and Roach, 2014). Overall, the extent to which academic and industrial logics translate into fundamentally different skill demands in practice remains unclear.

Answering my research questions requires estimating each individual's productivity in both Academia and Industry, that is, their counterfactual productivity across sectors. Doing so involves recovering two productivity parameters per individual, which goes beyond standard approaches that focus on average treatment effects. This is empirically challenging for several reasons. First, it requires detailed data about individuals' employment and their productivity. Second, because most individuals are only observed in one sector, I lack information on their productivity in the sector they did not join. Third, sorting is not random because individuals are expected to self-select into sector of employment based on earnings gains and preferences - most of which are unobserved - which prevents me from estimating counterfactual productivity using observed characteristics alone.

I address these challenges by using a confidential census of all STEM PhDs who graduated between 1970 and 2013 from an American university, with information about their sorting at graduation. To measure scientists' productivity, I merge this dataset with a confidential survey of scientists conducted during their careers, which provides information about their wages, as well as with publication data from Web of Science. I use information about scientists' wages as an empirical proxy for individuals' productivity (Agarwal and Ohyama, 2013; Campbell, Ganco, Franco and Agarwal, 2012; Carnahan et al., 2012), and publication records as a complementary measure of productivity in academic settings (Bikard et al., 2019; Ding et al., 2010; Levin and Stephan, 1991).

I recover each individual's academic and industrial productivity by estimating Marginal Treatment Effects (MTE) of sector choice on empirical proxies of productivity - namely, wages and publications.

This approach gives predicted productivity in both Academia and Industry for each individual, regardless of their actual sorting outcome. This method relies on an instrument to account for endogeneity arising from individuals' self-selection into sector. The first step estimates each individual's probability of entering Industry based on observables and the instrument. Among individuals with similar predicted probabilities, some are observed in Academia and others in Industry. This variation allows me to trace academic and industrial productivity across different values of the predicted probability of entering Industry and, in turn, to recover each individual's potential productivity in both sectors. My instrument leverages plausibly exogenous variations in firms' demand for field-specific scientific labor over time. Specifically, I use the year in which a PhD student graduates to generate variation in their likelihood of entering Industry, conditional on PhD major and macroeconomic conditions (Arellano-Bover, 2024; Oyer, 2006).

My first set of results examines the extent to which academic productivity reflects a scientist's productivity in an industrial setting. Using individuals' predicted earnings in Academia and in Industry, I find that the correlation between academic and industrial productivity is 0.28. This implies that academic productivity is only moderately informative about industrial productivity. For example, individuals with above-median productivity in Academia have a 62% probability of also being above-median productivity in Industry. On the other hand, individuals with below-median productivity in Academia have a 38% probability to have above-median productivity in Industry. Using predicted publications in Academia instead of predicted earnings as a proxy for academic productivity yields a correlation of 0.19, reinforcing the idea that academic productivity only partially reflects industrial productivity. Moreover, this correlation has decreased by 40% between 1975 and 2010, implying that academic productivity has become a less reliable predictor of industrial productivity, a pattern consistent with broader shifts in the relationship between academic science and industrial application (Arora et al., 2018).

Given the distribution of productivity among individuals, my second set of results examines how talent is allocated across sectors. Among those hired into Industry, 79% have above-median industrial productivity, while 21% fall below the median. These lower-productivity hires are significantly more likely to be laid off, suggesting a mismatch with the industrial setting. To explore whether such hires reflect evaluation constraints arising from the limited ability of academic productivity to proxy for industrial potential, I leverage variation across PhD majors in the alignment between academic and industrial productivity. I find that in majors where this alignment is weaker, firms' hires are more likely to have low productivity and to be laid-off, consistent with firms facing greater challenges in evaluating candidates. Moreover, in these majors, low-performing hires are more likely to display indicators of academic productivity at graduation, such as publications or degrees from top institutions. These patterns provide suggestive evidence of an inference challenge: when academic performance is a noisy proxy for industrial productivity, firms face greater difficulty identifying candidates who will

perform well in Industry, increasing the likelihood of costly mismatches.

Finally, I examine how firms can attract high-potential scientists. I begin by analyzing stated preferences for job attributes and find substantial heterogeneity, even among individuals with above-median industrial productivity. Those with high industrial productivity only place greater emphasis on advancement opportunities, responsibility and salary, while those with high industrial *and* academic productivity are more likely to prioritize intellectual challenge and independence. I then use observed sorting outcomes and predicted earnings potential to estimate the earnings premium required to induce scientists to choose Industry. The results reveal substantial variation in compensating differentials: 40% of scientists would require a premium equivalent to 3.8 times the academic wage to be indifferent between sectors, while 8% require no premium and in some cases, appear willing to accept lower earnings to work in Industry. Importantly, preferences are correlated with productivity: scientists with high industrial potential tend to have weaker preferences for Academia. As a result, attracting individuals with above-median industrial productivity requires an earnings premium between 0.2 and 1.1 times the academic wage. These findings underscore that more accurate evaluation of scientists can also help firms avoid offering high premiums to individuals unlikely to perform well in Industry.

This paper makes several contributions. First, it contributes to the strategic human capital literature by showing how institutional context shapes firms' ability to hire high-skilled talent (Bidwell, 2011; Bidwell and Mollick, 2015; Coff and Kryscynski, 2011). While most work has focused on workers' mobility and retention, we still know little about how firms hire under uncertainty about workers' skills, as well as the downstream consequences of these decisions for firm outcomes (Campbell, Ganco, Franco and Agarwal, 2012; Coff, 1997; Marx et al., 2009; Ng and Stuart, 2022; Sevcenko and Ethiraj, 2018; Starr et al., 2019, 2018). I show that when indicators of productivity originate from a different organizational environment, firms may face constraints in evaluating and attracting candidates, even when those indicators are perfectly informative about worker productivity in their original context. These constraints have tangible and costly implications for firms such as layoffs, especially in fields where cross-context productivity alignment is weak. Methodologically, I decouple actual sorting outcomes from potential productivity in both sectors, answering recent calls to consider counterfactual worker-organization pairings (Coff and Rickley, 2021). I do so by applying Marginal Treatment Effects (MTE) (Brinch et al., 2017; Heckman and Vytlacil, 2007), an econometric approach that remains uncommon in the strategy literature, but is uniquely suited to capture individual-level heterogeneity and recover counterfactual outcomes under selection.

Second, this paper contributes to the innovation literature on how knowledge workers - and scientific knowledge more broadly - shape firm performance (Baruffaldi and Poege, 2025; Gambardella et al., 2015; Palomeras and Melero, 2010; Polidoro Jr and Theeke, 2012; Shvadron, 2023; Singh and Agrawal, 2011). Prior work has often treated scientists as a relatively homogeneous input into firms' R&D processes (Arora et al., 2024, 2025) or focused on individuals' productivity *within* sectors (Azoulay

et al., 2017, 2011; Ganco, 2013; Ganco et al., 2015; Oettl, 2012). This leaves open questions about the portability of talent across institutional environments (Agrawal, 2001). I show that the value of scientific talent is both highly heterogeneous and dependent on the institutional context at stake. In particular, my results show that the distinct logics of Industry and Academia fundamentally shapes *which* individuals are productive *where* (Bikard, 2018; Conti and Visentin, 2015; Dasgupta and David, 1994; Sauermann and Roach, 2014; Stephan, 2012). In doing so, this paper contributes to the absorptive capacity literature (Cockburn and Henderson, 1998; Cohen and Levinthal, 1990) by showing that the ability to evaluate and integrate external knowledge also hinges on selecting the right individuals. Prior work has emphasized the importance of identifying and pursuing external opportunities that align with a firm's strategic goals (Arora and Gambardella, 1994; Tranchero, 2023). I show that this logic applies to hiring decisions as well: leveraging scientific talent is not simply about acquiring any highly trained individuals. Rather, it depends on firms' ability to identify and attract those whose skills align with the firm's goals and organizational context.

The paper proceeds as follows. Section 2 introduces the conceptual framework. Section 3 presents the estimation method. Section 4 discusses the empirical setting, data and measurement strategy. Section 5 describes the research design. Section 6 reports the results. Section 7 shows additional robustness analyses and Section 8 concludes.

2 Conceptual Framework

In this section, I develop a framework that models productivity as sector-specific, resulting from the interaction between individuals' skills and the organizational context. A central feature of this framework is that it decouples individuals' productivity in each sector from their sorting outcome. The framework proceeds in three parts. First, I draw on strategic human capital theory to examine the hiring challenges of selecting and attracting individuals when productivity is context-specific and unobserved ex-ante. Second, I show how firms' hiring challenges can be conceptualized through the joint distribution of sector-specific productivity, which captures the alignment (or misalignment) between individuals' productivity across settings. Third, I apply this framework to the case of Academia and Industry, and discuss why it is theoretically unclear whether academic productivity reflects scientists' potential in an industrial setting.

2.1 Hiring Under Uncertainty About Individuals' Productivity

Human capital is widely recognized as a key source of competitive advantage, particularly in knowledge-intensive sectors where individuals' technical expertise, creativity and domain-specific tacit knowledge drive organizational outcomes (Barney, 1991; Coff and Kryscynski, 2011; Coff, 1997). A central insight from this literature is that productivity is not solely a function of individuals'

skills, but also of how those skills interact with the organization itself, an idea often used to explain firm-specific human capital (Lazear, 2009). From this perspective, individuals represent bundles of skills accumulated through education, experience, and innate ability, whose value depends on the organizational context in which they are deployed (Becker, 1962). That is, the value of human capital is both individual-specific and context-dependent: the same person may be highly productive in one setting and generate less value in another. As a consequence, hiring the right individuals - i.e., those whose skills align with the organization's unique environment - is the foundational first step in building a valuable workforce that can generate competitive advantage (Kryscynski, 2021).

Hiring involves two related challenges: i) *selecting* promising candidates from the applicant pool and ii) *attracting* promising individuals into the applicant pool. First, firms must evaluate and select individuals whose skills will generate value under their own systems, cultures, and demands (Cappelli, 2019; Huang and Cappelli, 2010). Yet, at the point of hire, information about workers' abilities and fit is typically limited, especially in external hires (Bidwell, 2011; Coff and Kryscynski, 2011). What firms observe are often indicators, such as credentials, publications, or past roles, which originate from outside the organization. As a result, hiring becomes a problem of inference across contexts: firms must predict how someone will perform in their own environment using information produced in another. This suggests that evaluating workers across organizational boundaries may be difficult, especially when those settings reward different skills, values, or work styles. Second, firms must design offers that successfully attract high-potential individuals, without overpaying for individuals who might appear strong but are ultimately misaligned with the firm's needs (Prendergast, 1999; Roach and Sauermann, 2024). This too hinges on understanding the relationship between productivity across contexts, in order to identify whom to target and how much to offer.

2.2 Sector-Specific Productivity - A Simple Framework

Consider two sectors of employment, Academia and Industry, indexed by $j \in \{A, I\}$. Each sector j values a set of skills k^j. Individuals are indexed by i and possess a certain level of skills valued by each sector, k_i^j. This determines their productivity in each sector, P_i^j. Each individual is characterized by their two-dimensional vector: (P_i^A, P_i^I). Using these pairs of values, I define *'Academia-specialists'* as individuals who are primarily productive in Academia but not in Industry, *'Industry-specialists'* as individuals primarily productive in Industry but not in Academia, *'stars'* as individuals highly productive in both sectors and *'modest achievers'* as individuals *relatively* less productive than the others.[1] The higher P_i^j, the more productive individuals are in sector j. This means that the most valuable workers for firms are *'Industry-specialists'* and *'stars'*.

[1]Note that my definition of 'stars' differs from works in which 'stars' are defined with respect to their productivity *within* an organization or a sector (Agrawal et al., 2017; Carnahan et al., 2012; Hess and Rothaermel, 2011; Hsu and Kuhn, 2023; Liu et al., 2018; Oettl, 2012) irrespective of their productivity in other environments.

Take the perspective of Industry trying to hire. The joint distribution between P_i^A and P_i^I determines the extent to which individuals who are highly productive in Academia would also be highly productive in Industry. This relationship can be summarized by the correlation between P_i^A and P_i^I. Figure 1 illustrates three stylized cases. When the correlation between P_i^A and P_i^I is negative (Figure 1a), individuals tend to be either *'Academia-specialists'* or *'Industry-specialists'*. When the correlation is positive (Figure 1b), individuals tend to be either *'stars'* or *'modest achievers'*. When the correlation is zero (Figure 1c), individuals are present in the four quadrants. From Industry's perspective, this correlation is central to both the evaluation and attraction of talent. When the correlation is high, individuals' productivity in Academia provides reliable information about their productivity in Industry, enabling firms to make informed decisions about who to select and target. But when the correlation is low or negative, individuals' productivity in Academia becomes noisy or misleading, raising the risk that firms may select or attract low-performing individuals, or fail to hire high-productivity candidates.

2.3 The Institutional Logics of Industry and Academia

Theoretically, whether academic and industrial productivity are aligned - and if so, to what extent - is ambiguous. In Appendix C, I formally model the correlation between academic and industrial productivity. Here, I summarize the main intuition.

The correlation between academic and industrial productivity depends on two forces: the overlap in what each sector rewards (demand-side) and how skills are bundled across individuals (supply-side). The correlation will tend to be positive when: i) sectors primarily value a common foundation of skills or ii) when they primarily value different skills, but individuals tend to possess both set of skills. In contrast, the correlation will be negative when sectors primarily reward different skills, and individuals typically possess one set of skills but not the other. I now discuss successively the skills that Industry and Academia reward (demand-side alignment) and the structure of skill endowments across individuals (supply-side distribution).

Academia and Industry value common skills such as analytical reasoning, technical depth, and creativity, which support problem-solving and innovation in a range of settings (Sauermann and Roach, 2014; Stern, 2004). However, each sector is governed by its own institutional logic, which influences how productivity is defined and what skills are rewarded (Gittelman and Kogut, 2003; Perkmann et al., 2013; Tartari and Breschi, 2012). In Academia, the dominant logic emphasizes intellectual autonomy, peer recognition, and novel contributions to knowledge. Researchers are expected to generate novel ideas, work independently, and pursue theoretical insight (Hill and Stein, 2025; Merton, 1973). These incentives reward skills such as originality, self-direction, deep expertise in narrow domains, and the ability to articulate theoretical contributions. In contrast, Industry operates

under a logic of execution, collaboration, and market relevance. Knowledge workers must work under time and resource constraints, and their output must advance products, improve processes, or create commercial opportunities (Agarwal and Ohyama, 2013; Bikard, 2018; Sauermann and Stephan, 2013). This institutional context rewards skills such as cross-functional teamwork, strategic prioritization, responsiveness to feedback, and the ability to deliver usable outcomes on tight timelines. Recent work has also highlighted a structural shift: firms have moved away from upstream scientific research toward downstream development and application, narrowing the overlap in skills demanded across sectors (Arora et al., 2018). As firms reduce their engagement with basic science, the skills required for success in corporate R&D may become increasingly distinct from those fostered in academic environments. Overall, it remains unclear how much overlap exists in the skills each sector rewards.

On the supply-side, the correlation between academic and industrial productivity also depends on how skills are distributed across individuals. Even when sectors value distinct skill sets, academic and industrial productivity may still be positively correlated if individuals tend to possess both types of skills. Sector-specific skills need not be mutually exclusive: individuals may be naturally endowed with both sector-specific skills or develop them in parallel, through diverse experiences or intentional skill investments. However, when skill development is shaped by institutional incentives - for example, if academic environments discourage engagement with applied problem-solving - individuals may specialize, and productivity across sectors may be weakly or even negatively correlated. While direct evidence on skill co-occurrence is limited, prior work suggests that scientists often make strategic trade-offs based on what their environment rewards (Lazear, 2009; Sauermann and Roach, 2014; Stern, 2004).

In sum, the correlation between academic and industrial productivity is theoretically ambiguous. While both sectors value similar skills, it is unclear how much these contribute relative to sector-specific skills. Moreover, there is little evidence on whether sector-specific skills co-occur or trade off across individuals. Empirically estimating the joint distribution between academic and industrial productivity is therefore essential to examine the challenges firms may face in hiring scientific talent.

3 Empirical Framework

3.1 Parameters of Interest and Empirical Challenges

To estimate P_i^A and P_i^I as per Figure 1, I must recover two productivity parameters for each individual. There are two main challenges in doing so. First, individuals join one of the two sectors, so I only observe their productivity in Academia or in Industry.[2] Second, I cannot simply recover an

[2]Some individuals might be changing sector of employment during their career, allowing me to observe their productivity in both sectors. However, this is a limited and selected sample.

individual's counterfactual productivity in the sector that they did not join by making inference based on observable characteristics. This is because we expect individuals to self-select into sector of employment based on unobserved characteristics (e.g., preferences) which might be correlated with their productivity. For example, individuals who self-select into Industry may also be more productive there than observationally similar individuals who remain in Academia, leading to a classic problem of selection on unobservables.

3.2 Estimation Technique

One standard approach to account for endogeneity arising from unobservable components is to estimate Two-Stage Least Squares (TSLS) using an instrumental variable. However, this entails significant limitations when trying to estimate separately P_i^A and P_i^I: (i) TSLS estimates a *difference* between two potential outcomes (i.e., $P_i^B - P_i^I$), rather than the two potential outcomes separately, and (ii) TSLS identifies this effect only for a local subset of individuals - the compliers - rather than across the full distribution. To recover P_i^A and P_i^I separately for each individual, I estimate Marginal Treatment Effects (MTE) (Björklund and Moffitt, 1987; Cornelissen et al., 2016; Heckman and Vytlacil, 1999, 2007). MTE also relies on the use of an instrument, but provides a more general framework than TSLS. I present here the main intuition behind the MTE estimation and refer the reader to Appendix D for a more comprehensive discussion.

MTE relies on the same key ingredients as TSLS: i) an outcome (in this paper, productivity) ii) an independent variable (joining Industry at graduation)[3] and iii) an instrument. The first step of the MTE estimation is similar to the first stage of a TSLS: I estimate the probability that individual i enters Industry as a function of observed covariates X_i and an instrument Z_i. This yields a propensity score, $P(Z_i, X_i)$: the predicted probability that individual i joins Industry given their values of X_i and Z_i.

The key insight of the MTE approach is that, conditional on the same propensity score, some individuals are observed in Industry and others in Academia. This is because individuals differ in unobserved ways (e.g., their preferences or comparative advantage) which influence sector choice even among those with the same predicted likelihood of entering Industry based on X_i and Z_i. These unobserved factors, denoted by U_{D_i}, are assumed to be uncorrelated with the instrument itself.[4] For each value of the propensity score $P(Z_i, X_i) = p$, individuals observed in Academia must have $U_{D_i} \geq p$ and their average outcome informs us about $E[P_i^A | X_i = x, U_{D_i} \geq p]$. Similarly, individuals observed in Industry must have have $U_{D_i} \leq p$, and their average outcome informs us about $E[P_i^I | X_i = x, U_{D_i} \leq p]$. By modeling observed outcomes as smooth functions of the propensity

[3] I use sector joined *at graduation* because it allows me to estimate productivity trajectories from the beginning of individuals' careers. However, the results have broader implications for how firms evaluate and attract talent throughout the career, including through mid-career hiring or poaching.

[4] U_{D_i} is often called 'unobserved resistance' or distaste for treatment.

score within each sector, we can estimate the expected potential outcome at each value of U_D, i.e., $E[P_i^A|X_i = x, U_{D_i} = u]$ and $E[P_i^I|X_i = x, U_{D_i} = u]$. Using individuals' observed characteristics and realized sector choice D_i, I can infer their location in the unobserved heterogeneity distribution. This allows me to generate corresponding predictions for individuals' productivity in both sectors, i.e., $E[P_i^A|X_i, D_i, p_i]$ and $E[P_i^I|X_i, D_i, p_i]$.

The assumptions needed for MTE are the same as for TSLS: relevance, exclusion and monotonicity. In theory, it is possible to estimate MTE with no further assumption. However, this requires full support of the propensity score for both those joining Academia and those joining Industry for all values of X, which is rarely feasible in practice (Andresen, 2018). Following the literature, I impose a functional form assumption of additive separability between observable and non-observable components which enables identification of the MTE over the common support unconditional on X (Carneiro et al., 2011; Cornelissen et al., 2016).

4 Setting, Data and Measurement

4.1 Data

I use individual-level data from three confidential files: (i) the Survey of Earned Doctorates (SED), (ii) the Surveys of Doctorate Recipients (SDR) and (iii) the 2015 Survey of Doctorate Recipients Bibliometric Research Data (SDR15-WoS). These surveys are conducted by the National Center for Science and Engineering Statistics (NCSES) and cover individuals who received a PhD from an American institution.[5]

The SED is an *annual census* of all individuals who graduated from a PhD in a given academic year, conducted since 1958 (e.g., the SED 2010 is sent to all individuals who earned their doctoral degree between 1 July 2009 and 30 June 2010).[6] The response rate is about 90%, providing almost full coverage of the population of individuals who receive their PhD in the United States in a given year. The SED contains information about doctoral studies (e.g., PhD field of study and institution) as well as information about previous education (e.g., bachelor and master). It also contains demographics and background information (e.g., place of birth, gender, race, parental education, citizenship status). Importantly for the analysis, the SED contains information about postgraduation plans (e.g., work, postdoc, other study or training).

The SDR is a *biennal panel survey* of doctorate holders conducted since 1973 which contains longitudinal data about individuals' employment at the time of survey - i.e., during their career - such

[5]These data are restricted-use and were accessed remotely from the National Opinion Research Center (NORC) data enclave.

[6]For reference, about 55k research doctorates were awarded in the United States in 2020. This excludes recipients of professional doctoral degrees, such as MD, DDS, DVM, JD, DPharm, DMin, and PsyD.

as sector of employment, principal job activity and earnings. For each survey, the sample consists of both individuals who were surveyed in previous SDR editions (as long as they are less than 76 years old) as well as new doctoral graduates who earned their PhD since the last cycle. The response rate is about 65%. Individuals can be surveyed several times during their career and hence might appear in several SDR waves. While the SDR reports the institution name for individuals who work in the academic sector, I am not allowed to access employer name for those who don't.

In the years 2022/2023, the NCSES launched a new Research Data Infrastructure with the goal of creating new linkages with external data. I take advantage of this effort by using a newly created dataset that matches each individual surveyed in the SDR 2015 to their research output from graduation until 2017, using information from Web of Science. This gives me individual-level information about publication output produced between graduation and 2017 such as publications, top publications and average number of citations in the first 2 or 5 years after publication. For each of these measures and for each individual, I have access to aggregate values up to 2017 as well as more granular information by bins of experience (e.g., publications produced during the PhD, publications produced during the first 5 years post-graduation and so on).[7] Note that I do not have information at the publication-level such as journal or citations.

I first match the SDR 2015 to the SDR15-WoS and the SED. This gives me information about all individuals surveyed during their career in 2015 regarding their background, doctoral education, employment plans at graduation and employment status in 2015, as well as publication output. I then match this sample to the other SDR surveys I have access to (survey years 1995, 1997, 1999, 2001, 2003, 2006, 2008, 2010, 2013, 2017 and 2019) to retrieve all the information available about individuals' earnings during their career.

The initial sample contains 78,320 individuals. I restrict the analysis to STEM fields and exclude individuals with a PhD in Humanities, Education, Business and Health due to limited observations. I retain cohorts graduating in or after 1970 because of several changes in the way the SDR answers are coded before that year. To assign individuals to a sector of employment at the start of their career, I keep students with definite postgraduate commitments at graduation at the time of survey (71%)[8] and exclude individuals who declare starting an internship, traineeship, clinical residency or military service (3%). I also exclude individuals who do not start their career in the United States (12%). Finally, because my estimation requires observing individuals' productivity at least once in their initial sector of employment, I drop individuals who are never observed working in their starting sector (18%). I discuss the robustness of this last restriction in Section 7.

[7] I have counts of these measures for the 5 years before graduation and for bins of 5 years after graduation, starting for bin of years 1 to 5 after graduation up to bin of years 46 to 50 after graduation.

[8] 9% are negotiating with one or more organizations, 18% are seeking a position but have no specific prospects, 1% are joining another full-time degree and 1% declare 'other' plans.

This represents 22,609 unique individuals. Because the earnings outcome is a flow and doctorate holders can be surveyed several times during their career, my sample contains 69,917 observations in total, meaning that I observe on average 3 earnings values per individual.

I refer to PhD *major* as the most granular field of study I have access to for each individual (e.g., Biology, Genetics, Biochemistry, Astronomy, Computer Engineering etc.). This is what students choose to study. My dataset includes 58 distinct PhD majors. I classify these majors into 5 PhD *fields*: Computer Sciences and Mathematics, Life Sciences, Physical Sciences, Social Sciences and Engineering (see Figure A.1 for a list of PhD fields and majors).

4.2 Dependent Variables: Measurement of Productivity

From a theoretical perspective, the ideal measure of an individual's productivity is their marginal product of labor, i.e., the value they create for their organization. In practice, this is rarely observable outside of narrowly defined production settings. The most widely used empirical proxy for productivity is earnings (Campbell, Ganco, Franco and Agarwal, 2012; Rosen, 1981). I also incorporate data on publication output as another measure of academic productivity. However, note that earnings has the advantage to capture overall productivity across all tasks - including teaching, administration, mentoring, and service - while publication output is more closely linked to *research-specific* productivity.[9]

Earnings: Earnings correspond to the total earned income before deductions in the previous year. This includes wages, bonuses, overtime, consulting fees and summertime teaching or research, and I have access to several earnings observations per individual at different years of experience.[10] I transform all values in 2015 US dollars. I winsorize observations below the 1^{st} percentile and above the 99^{th} percentile.

Publications: I calculate the number of publications published after PhD graduation by subtracting the number of publications published before graduation from the total number of publications published by the individual until 2017.[11] I winsorize observations below the 1^{st} percentile and above the 99^{th}

[9] One concern is that earnings in Academia are more likely to be influenced by institutional pay scales and tenure-based salary progression, which may compress wages by limiting variation across individuals. While this wage compression could attenuate measured differences in productivity, it is unlikely to significantly distort the rank-ordering of individuals. Moreover, academic earnings still correlate positively with other measures of productivity, such as publication counts, citation impact, and research funding (Hilmer et al., 2015; Katz, 1973).

[10] To avoid contaminating the estimation of P_i^A and P_i^I, I exclude earnings observations that occur after individuals switch sectors - that is, earnings corresponding to the sector not initially chosen. The MTE framework identifies potential outcomes by comparing individuals with similar propensity score who sorted into different sectors. Including post-switching earnings observations from individuals who switch sectors would distort this estimation. For example, using earnings in Industry from someone who initially joined Academia would improperly enter into the estimation of P_i^A. I show in Section 7 that results remain similar when including all earnings observations.

[11] I only include publications that were articles, reviews or conference proceedings.

percentile. Depending on the specification, I also use publications weighted with citations received in the first 5 years of publication. While publication output is widely used as a measure of scientific productivity, its interpretation varies by sector. In Academia, it arguably reflects performance incentives and norms. In Industry, however, publication behavior may be shaped by organizational policies, intellectual property concerns, or strategic disclosure practices, rather than underlying research productivity. As such, I use publications as a complementary measure of academic productivity, but rely primarily on earnings to proxy for productivity in industry.

4.3 Independent Variable: Measurement of Sector Joined

My main independent variable of interest is the sector that individuals join upon graduation. I first assign each individual to three potential sectors: Industry (business for-profit and self-employed), Academia (2-year and 4-year college or university and university-affiliated research institute) or Government (US or foreign government). My preferred treatment variable is an indicator variable equal to 1 for individuals joining Industry and 0 for individuals joining any other sector, i.e., Academia or Government (in the rest of the paper, I will loosely refer to this treatment variable as joining Industry vs Academia).[12] Overall, among the 22,609 unique individuals in my sample, 21% are joining Industry.

Note that 46% of individuals in my sample start their career in a postdoc position. For my main results, I choose to assign individuals doing a postdoc to the academic sector for two main reasons. First, if postdoc is considered as a first post-PhD employment, I should in theory assign individuals to the sector where the postdoc is realized. While I do not have such information for individuals who graduated between 1969 and 2003,[13] only 2% of individuals who graduated in or after 2004 did so in the private sector. If postdoc is rather considered as a continuation of PhD education, I should assign individuals to the first sector they join after their postdoctoral studies. Unfortunately, while I observe sector of employment for all individuals in my sample after PhD graduation, I cannot easily identify the sector where individuals work right after finishing their postdoctoral studies because (i) I do not know exactly when they finish their postdoc and (ii) I can't easily proxy for the end of postdoctoral studies because I do not necessarily observe individuals again in the early-part of the career, especially for older cohorts. Nevertheless, I show in Section 7 that my results are robust to discarding the first 10 years post-PhD graduation of postdocs and using the first sector of employment where postdocs are observed after these 10 years to define their sector of employment post-graduation.[14]

[12] I show in Section 7 that results are robust to excluding the Government category.

[13] The question started to be asked to individuals in 2004 onwards.

[14] I use 10 years as my threshold because Figure A.2 shows that the earnings curve of postdocs converges with the one of individuals who start in Academia without a postdoc at around 10 years of experience.

4.4 Descriptive Statistics

Figure A.3 shows the number of PhDs by cohort of graduation and PhD field in my sample. The number of individuals increases with cohort of graduation, reflecting the increase in the number of PhDs over time. Figure A.4 shows that about 20% of PhDs join the private sector upon graduation. There is, however, substantial heterogeneity across fields as showed in Figure 2: about 50% of PhDs in Engineering join the private sector, vs 10% in Life Science or Social Sciences.

Table 1 reports summary statistics at the individual level, by sector joined. Individuals who join the private sector earn on average across the career 67k more per year than individuals who started their career in Academia. In a simple regression of logged earnings on sector of employment, controlling for observable characteristics, Table B.1 shows that the earnings premium in Industry is about 40%. Individuals who join Industry also publish on average 16 fewer publications between graduation and the time of survey. In Academia, 44% of individuals are female, vs only 29% in Industry. Overall, 72% of individuals in my sample are Americans but a higher share of international students tend to join the private sector at graduation than Academia.[15]

Figure A.5 shows earnings as a function of career experience and sector joined. The earnings curve in Industry has a higher intercept but is less steep, so that earnings in both sectors look roughly similar in the raw data towards the end of the career (Levin and Stephan, 1991). This pattern could be due to differences in skill obsolescence (Deming and Noray, 2018), differences in ability and preferences (Roach and Sauermann, 2010; Stern, 2004) and differences in physical capital investments and complementarities between basic an applied scientists in Industry vs Academia (Agarwal and Ohyama, 2013). Figure A.6 shows citation-weighted publications as a function of sector joined and experience. The curve is upward sloping in Academia and noisier in Industry, likely because individuals' activity in Industry is less homogeneous, with workers i) not necessarily employed in research and/or ii) working in firms that do not allow them to publish research results.

5 Research Design

To estimate individual i's productivity in Academia, P_i^A, and individual i's productivity in Industry, P_i^I, I estimate Marginal Treatment Effects as presented in Section 3. This method requires the use of a variable to instrument for sector joined at graduation. This is because sorting is jointly determined by firms' and universities' demand, as well as individuals' self-selection. My strategy aims to leverage

[15] I follow the classification of the SED and identify U.S. citizens as individuals who are either U.S. native born (68% of my sample), U.S. naturalized (4% of my sample), U.S. unspecified (not present in my sample) or U.S. citizen assumed for Puerto Ricans who did not report citizenship status (not present in my sample). Individuals who are not identified as U.S. citizens can be non-U.S. immigrant permanent resident (5%), non-U.S. non-immigrant temporary resident (23%) or non-U.S. visa status unknown (0.1%).

variation in labor demand from firms versus universities at the time individuals graduate from their PhDs.

Firms' demand for PhDs fluctuates across years due to factors such as technological advancements and shifts in R&D investment priorities (Toole and Czarnitzki, 2010). For instance, breakthroughs in gene-editing technologies like CRISPR can lead to a surge in demand for PhDs in genetics and biotechnology, as firms seek to explore new applications in agriculture and medicine. Similarly, universities often adjust their hiring based on evolving hiring needs, academic priorities, funding availability, and strategic research initiatives (Goolsbee and Syverson, 2023). As a result, PhD students in the same major may face very different employment landscapes depending on the year they graduate. Some cohorts face stronger relative demand from firms, others from universities. I use this cohort-level variation in relative labor demand to construct my instrument, exploiting the fact that timing of graduation shapes the likelihood of entering Industry vs Academia. This approach is similar in spirit to Oyer (2006) and more recently to Arellano-Bover (2024).

Empirically, the instrument can be thought of as a vector of PhD major $\times$ PhD cohort fixed-effects, controlling for PhD major fixed-effects.[16] In practice, I create a leave-one out continuous equivalent to reduce the risk of bias (Dobbie et al., 2018; Sampat and Williams, 2019). For each individual, the instrument Z_i is therefore equal to:

$$Z_i = \frac{\sum_{i'} Industry_{i'} \times \mathbf{1}\{m(i') = m(i)\} \times \mathbf{1}\{c(i') = c(i)\} - Industry_i}{\sum_{i'} \mathbf{1}\{m(i') = m(i)\} \times \mathbf{1}\{c(i') = c(i)\} - 1} \tag{1}$$

with i' indexing individuals, $Industry_i$ an indicator equal to 1 if individual i joins Industry and 0 otherwise, $m(i)$ individual i's PhD major and $c(i)$ individual i's PhD cohort of graduation.

To be valid, the instrument: (i) should not impact the outcomes directly (i.e., other than through its impact on sector joined at graduation) and (ii) should not be correlated with unobserved factors that also influence individuals' sector choice, such as individuals' skills or their preferences. Figure A.7 formalizes how these assumptions relate to key variables in a simple model. I discuss both assumptions successively considering the earnings outcome.[17]

Regarding (i), I need to control for the overall level of demand in order to only use as variation the *relative* demand of firms vs universities. Said differently, I want to control for general macroeconomic conditions that would be correlated with both Z_i and earnings. For instance, if I observe a large share of individuals starting in Industry for a specific major/cohort at a time where the economy is also at its best, earnings in the private sector might appear high even though individuals would have enjoyed high earnings in Academia too. To mitigate this concern, I include a control for the unemployment rate

[16] I follow the literature and only keep PhD major $\times$ PhD cohort fixed-effects for which I have at least 10 observations (Heckman, 1981).

[17] The discussion remains similar considering publications.

at the time of PhD graduation which aims to control for general macroeconomic conditions when the individual enters the labor force.[18] I also control for field-specific changes in funding or opportunities that would impact both sectors by controlling for field-specific federal funding the year of graduation.[19]

For assumption (ii) to hold, there should be no unobservable component that is correlated with both the instrument and individuals' sorting. In particular, this implies that individuals who are more productive in Industry (or said differently, who would be paid more in Industry) do not systematically graduate in years where firms' vs universities' demand is high. This assumption would break in the following cases: strategic graduation ; strategic entry ; skill composition shifts.

First, one concern could be that individuals who would be paid more in Industry strategically adjust their PhD graduation year based on firms' demand at graduation. For instance, if firms' demand is low at the time of graduation, these individuals could decide to defer their graduation hoping that demand will be higher the year after. Similarly, high-potential earners in Industry who should have graduated the following year might rush to graduate to seize favorable opportunities. To test for this threat, I use information about PhD entry year to calculate PhD duration and regress this variable on the instrument.[20] Table 2 Column (1) shows no significant correlation between the time that individuals take to complete their PhD and firms' vs universities' demand. To account for potential offsetting behaviors (some individuals extending, others rushing), I also run two similar regressions where the outcomes are: an indicator equal to 1 if individuals' PhD duration is strictly below the average PhD duration of individuals with the same PhD major and PhD entry year combination and an indicator equal to 1 if individuals' PhD duration is strictly above the average PhD duration of individuals with the same PhD major and PhD entry year combination.[21] Columns (2) and (3) show no significant correlation. Overall, it does not seem that individuals manipulate their graduation year in a way that would be correlated with my instrument.

Another concern could be that high potential earners in Industry are able to forecast Industry demand in following years and strategically time their PhD entry. I run two additional tests to examine this threat. First, I regress the value of the instrument at graduation on the value of the instrument the year before PhD entry.[22] Table 2 Column (4) shows that there is no significant correlation, suggesting that conditions at graduation are not correlated with conditions that individuals would observe at the time they apply for a PhD. I also check whether market conditions the year before entry predict future sorting. Column (5) shows no significant effect.

Finally, gradual shifts in skill composition across PhD cohorts could correlate with changes in firms'

[18] Source: `https://fred.stlouisfed.org/series/UNRATE`

[19] Source: `https://www.aaas.org/programs/r-d-budget-and-policy/historical-trends-federal-rd`

[20] PhD duration is calculated as the difference between PhD graduation year and PhD entry year. The number of observations is slightly lower than the overall sample because some individuals did not report PhD entry year.

[21] The average PhD duration year for each individual in a field/cohort is calculated excluding the focal individual.

[22] This is because individuals apply to PhD programs usually the year before starting their program.

versus universities' demand. Said differently, even if individuals do not strategically time their PhD graduation or entry based on anticipated market conditions, broader structural trends may lead to changes in who enters particular fields over time. For example, if Industry demand in certain fields increases steadily, individuals with stronger Industry-oriented skills may be disproportionately drawn to those fields, altering the average skill composition of graduating cohorts. This concern is most relevant in Computer Science, where Industry placement rises towards the end of my sample - though my dataset ends before the 2013 AlexNet breakthrough which likely accelerated this trend. I check in Section 7 that my results are robust to the inclusion of five-year-window fixed-effects interacted with PhD field. This allows me to compare individuals graduating within narrow time bands and mitigate the risk of long-run shifts in skill composition.

Finally, another worry is that students from specific institutions have systematically different career prospects or preferences. These differences might lead them to respond more or less strongly to firms' demand. If these students also tend to have systematically higher earnings, failing to control for doctoral institution could bias the estimates by conflating the effect of firms' demand with institution-specific factors that influence both sector choice and earnings. To mitigate this concern, I control for doctoral institution with a vector of fixed-effects. In a similar fashion, I control for gender, race, birthplace, citizenship and parental education.

Overall, any remaining threat to identification would require unobserved, time-varying shifts in the composition of PhD graduates within a major that are correlated with fluctuations in relative labor demand and not captured by my existing controls or addressed by my previous tests.

6 Results

6.1 First-Stage

I run the following first-stage regression at the individual level:

$$Industry_i = \gamma_0 + \gamma_1 Z_i + \gamma_2 \mathbf{X_i} + \nu_i \tag{2}$$

where $Industry_i$ is an indicator equal to 1 if individual i joins Industry and 0 otherwise. Z_i is the instrument for individual i, which is calculated at the PhD major $\times$ PhD cohort level. Following my identification strategy, $\mathbf{X_i}$ includes PhD major fixed-effects so that I rely on variation *across years within PhD major*, as well as controls for the unemployment rate and field-specific federal funding at the time of graduation. $\mathbf{X_i}$ also includes a linear and quadratic terms to control for experience as I observe individuals at different points in time during their career. It also includes an indicator equal to 1 if the individual declares being female and 0 otherwise, an indicator equal to 1 if the individual

declares being white and 0 otherwise, an indicator equal to 1 if at least one of the individual's parents has a doctoral degree and 0 otherwise, an indicator for being an American citizen and 0 otherwise, birth place fixed-effects and doctoral institution fixed-effects.[23] Standard errors are clustered at the PhD major $\times$ PhD cohort level since this is the level of variation of my instrument.

Figure A.8 shows an histogram of the instrument Z_i which varies between 0 and 0.93 with a mean of 0.20. Table B.2 shows the first-stage regression: a 10 percentage point increase in firms' vs universities' demand increases the likelihood of joining the private sector by 3%, with a Kleibergen-Paap F-statistic equals to 39. Figure A.9 shows histograms of the propensity score (predicted probabilities of joining Industry as a function of observables and the instrument), separately for individuals who join Industry and individuals who join Academia. There is large overlapping support for the two groups, which helps the MTE estimation by ensuring that individuals from both sectors are observed across a wide range of predicted probabilities. The distribution of the propensity score is shifted to the right for individuals who start in Industry, reflecting their higher likelihood of joining the private sector.[24]

6.2 Relationship Between Academic and Industrial Productivity

Figure A.10 presents the distributions of individuals' predicted *earnings in Academia*, which proxy for their productivity in Academia, and individuals' *predicted earnings in Industry*, which proxy for their productivity in Industry. I estimate these using a first-degree polynomial in the MTE specification and show in Section 7 that my results are robust to other specifications. My estimations were performed using the log of earnings, but values showed in figure A.10 are transformed into thousands of dollars for ease of interpretation.[25] As expected, potential earnings in Industry are generally higher than in Academia (118k on average vs 88k) and have higher variance (standard deviation of 44k vs 25k).

To what extent does academic productivity reflect a scientist's potential in an industrial setting? To answer this question, I first residualize each individual's predicted earnings in Industry and Academia by experience and its square. This step accounts for systematic earnings growth over the life cycle, allowing for more meaningful comparisons across individuals at different career stages. I then average these residualized values across observed periods to obtain a single productivity estimate per sector per individual. This yields a pair of productivity measures for each person, one for Academia and one for Industry. I then plot individuals' industrial productivity against their academic productivity. Results are showed in Figure 3a. I find a positive correlation between academic and industrial productivity equal to 0.28. This indicates moderate alignment between academic and industrial productivity: while academic productivity provides some information about a scientist's potential in Industry, it is a

[23] Birth place is at the state level for individuals who are born in the U.S. and at the country level for individual who are born outside of the U.S.

[24] I trim the thinnest tails of support (0.05% of the sample). This is to avoid extrapolation with very few observations and improve robustness of the estimates. This represents 34 observations. Results are largely insensitive to this trimming.

[25] Figure A.11 shows the histograms with potential earnings expressed in log.

relatively imprecise indicator. This result implies that an individual with above-median productivity in Academia has a 62% probability to also have above-median productivity in Industry. Conversely, an individual with below-median productivity in Academia has a 38% probability to have above-median productivity in Industry.

To complement this result, I also proxy for academic productivity using publication output. Using the same MTE approach previously applied to earnings, I estimate each individual's potential publication count in Academia, weighted by citations. To make this measure comparable across individuals, I residualize predicted publication counts by experience and field, given well-documented variation in publication norms across disciplines. This provides a standardized, productivity-based measure of research output over the career. Figure 3b shows the result. I find that the correlation between potential publications in Academia and potential earnings in Industry is 0.19, confirming that research output offers limited insight into industrial potential.

Looking at heterogeneity across observables, I find that the correlation between potential earnings is lower for females than for males (0.15 vs 0.25) and lower for Americans than for internationals (0.30 vs 0.36). This correlation is also stronger in Social Sciences (0.51) and in Computer Sciences (0.34) than in Life Sciences (0.18), Physical Sciences (0.13) and Engineering (0.11). This last result indicates that academic productivity in Engineering offers little information about an individual's potential in Industry.

Finally, I examine how the correlation between academic and industrial productivity has changed over time. I calculate the correlation between residualized potential earnings separately for each PhD cohort. Figure 4 shows the result. This correlation has decreased from 0.39 for cohorts before 1980 to 0.23 for cohort after 2010, representing a 40% decline. This implies that academic productivity has become a less reliable predictor of industrial productivity. This pattern is consistent with broader shifts in the relationship between academic science and industrial application, such as firms increasingly disengaging from Science (Arora et al., 2018).

6.3 Sorting and Realized Productivity of Firms' Hires

Building on the joint distribution between academic and industrial productivity, I now examine which individuals join the private sector. While I have estimated individuals' academic and industrial productivity independent of the sector they actually join, note that final sorting patterns reflect equilibrium outcomes involving both where individuals apply and who firms choose. I first use sorting outcomes to shed light on the constraints firms face in *evaluating* scientists using academic productivity. In Section 6.4, I shift focus to the supply-side and study scientists' preferences, quantifying the earnings premiums required to induce them to work in Industry and thus shedding light on how firms can *attract* high-potential candidates.

To implement this, I discretize residualized potential earnings in Academia and Industry into deciles and construct a 10×10 matrix of potential earnings combinations. Figure 5 displays this matrix.[26] Each cell represents a group of individuals who fall into a given decile of productivity in Industry and in Academia. The color of each cell indicates the sector individuals join: yellow denotes a majority entering Industry, dark purple indicates a majority entering Academia, and pink reflects a roughly even split between the two sectors.

Among individuals hired by firms, 79% have above-median industrial productivity (35% 'stars' and 44% 'Industry-specialists') while 21% have below-median industrial productivity (3% 'Academia-specialists' and 18% 'modest achievers'). In contrast, the composition of individuals hired by universities is more varied: 52% have above-median academic productivity (29% 'stars' and 23% 'Academia-specialists') and 48% have below-median academic productivity (12% 'Industry-specialists' and 36% 'modest achievers').

Figure 6 presents results by PhD field. In Computer Science, 80% of firms' hires have above-median industrial productivity, with a nearly even split between 'stars' and 'Industry-specialists'. In the Life Sciences and Physical Sciences, about 60% of firms' hires have above-median industrial productivity, but these individuals are primarily concentrated in the pool of 'Industry-specialists'. In Engineering, 55% of firms' hires have above-median productivity in Industry (18% 'stars' and 37% of 'Industry-specialists'). Figures A.15 and A.16 show that firms' hiring composition is broadly similar across gender and nationality: in each group, about 79% of hires have above-median productivity in Industry.

The previous result indicates that, among firms' hires, 21% have below-median industrial productivity. To assess whether firms seek to hire individuals who will be productive in Industry, I examine whether lower-productivity hires are more likely to be laid off. Each time an individual is surveyed, I observe whether they are currently working or not, and if not, the reported reason for unemployment. I also observe whether the individual remained with the same employer over the past two years and, if not, the reported reason for the change. Based on this information, I define a layoff indicator at the person-year level. It is equal to one if the individual either (i) reports currently not working due to a layoff, or (ii) reports having changed employer in the past two years due to a layoff. I then aggregate to the individual level, assigning a value of one if the person was ever laid off across any of the survey waves in which they appear. This variable aims to capture whether the individual ever experienced a layoff during the observation window, serving as a proxy for poor matches from the firm's perspective.

I then regress this outcome on a set of dummies for productivity group ('star', 'Industry-specialist', 'Academia-specialist', 'modest achiever'). I include demographic controls and PhD major fixed-effects. Since individuals who are surveyed more frequently have a mechanically higher chance of having been laid-off, I also add PhD cohort fixed-effects and a set of fixed-effects for the number of survey waves

[26] Figures A.13 and A.14 show the result in level.

in which each individual is observed. To account for differences in layoff risk across career stage, I additionally control for the experience level at which I first observe each individual. The results, reported in Table B.3, show that individuals with below-median industrial productivity ('Academia-specialists' and 'modest achievers') are significantly more likely to be laid off than those with higher predicted industrial productivity. This is consistent with the idea that lower-productivity hires are more likely to reflect firm-side mismatches, i.e., cases where individuals did not meet productivity expectations set at the time of hire.

To shed light on whether hiring mismatches stem from evaluation constraints, I exploit variation in the correlation between academic and industrial productivity across PhD majors.[27] For each major, I calculate the share of firms' hires who were ever laid off and regress this outcome on the major-level correlation. I then repeat the analysis using an alternative measure of mismatch: the share of firms' hires within each major whose predicted industrial productivity falls below the median.[28] Table 3 shows a negative and statistically significant relationship in both cases: in majors with weaker alignment between academic and industrial productivity firms exhibit higher rates of both layoffs and below-median industrial productivity hires. As expected, the estimate is attenuated in the layoff specification, consistent with noise in the measurement of layoff status.

I also examine whether individuals identified as mismatches are more likely to display strong academic indicators at graduation in majors where academic productivity is less predictive of industrial productivity. The broader idea behind this test is to assess whether firms are acting on the observable information they have at the time of hire, and whether reliance on academic attributes is more likely to lead to suboptimal hiring outcomes when academic productivity is a noisy predictor of industrial productivity. In other words, when academic and industrial productivity are weakly aligned, firms may rationally hire individuals who appear strong on paper but ultimately underperform in Industry. I consider three indicators of academic productivity at graduation: (i) publication count at graduation (ii) a binary indicator for having at least one publication at graduation,[29] and (iii) an indicator for having graduated from a top-20 U.S. PhD.[30]

I then estimate individual-level regressions among firms' hires who are either laid-off or have below-median industrial productivity, regressing indicators of academic productivity at graduation on

[27] This correlation varies between -0.10 and 0.55 with a mean of 0.18.

[28] The layoff-based measure reflects whether the match was ultimately dissolved by firms and thus provides a revealed-performance outcome. However, layoffs are relatively rare, especially as I only observe individuals on average three times during their career. The below-median productivity measure is available for all hires and provides a consistent proxy for performance quality, but does not indicate whether the match ended. Taken together, the two measures offer complementary insights into mismatch patterns.

[29] Measures (i) and (ii) are standardized to account for major-level differences in publication rates. On average, students have 2.2 publications at the time of graduation. 44% of students have no publication at the time of graduation.

[30] `https://www.timeshighereducation.com/student/best-universities/best-universities-united-states` US Rank 2025.

the correlation between academic and industrial productivity in their PhD major. All specifications control for demographic characteristics and PhD cohort fixed-effects. Table 4 reports the results. Both publication-based measures are negatively and significantly associated with the major-level correlation across both mismatch definitions. The top-institution indicator is also negatively associated with the correlation, but reaches significance only in the below-median productivity sample.[31] Overall, mismatched hires are more likely to display strong indicators of academic productivity at graduation in majors where academic and industrial productivity are weakly aligned.

Taken together, this analysis sheds light on evaluation constraints in hiring and their consequences for firms. I show that 21% of firms' hires have below-median industrial productivity and that these individuals are significantly more likely to be laid off, suggesting that some hires fail to deliver sustained value. To understand why such outcomes arise, I exploit variation in the correlation between academic and industrial productivity across PhD majors. Low-productivity hires and layoffs are systematically more common in majors where this alignment is weaker. Moreover, in those majors, these individuals are more likely to display strong indicators of academic productivity at graduation. This pattern provides evidence of an evaluation constraint, where measures of productivity developed in a different institutional environment can be noisy predictors of industrial productivity. These constraints have concrete consequences for firms including layoff costs, along with associated productivity losses and missed opportunities to hire better-matched candidates.

6.4 Wage Premiums and Talent Attraction

So far, I have used sorting outcomes to examine the productivity of firms' hires and shed light on evaluation constraints at the point of hire. In this section, I turn to the other side of the sorting process: the incentives that shape where scientists choose to work. Specifically, I use observed sectoral choices and predicted earnings in each sector to examine how much firms must offer to induce scientists to select into Industry, i.e., compensating differentials. This exercise helps quantify how much additional compensation is required for firms to attract high-potential individuals, highlighting not just informational constraints, but also the challenges firms face in making Industry an appealing option for top scientific talent.

I start by qualitatively examining the preferences of each pool of talent. I use questions asked to individuals about several job attributes: (1) opportunities for advancement (2) benefits (3) intellectual challenge (4) degree of independence (5) location (6) level of responsibility (7) security (8) salary and (9) contribution to society. For each attribute, individuals are asked whether they consider it

[31] As a robustness check, I use predicted academic productivity (from the MTE model) as the outcome. Results are similar, suggesting that the relationship holds when using a more comprehensive, model-based measure of academic performance.

as 'very important', 'somewhat important', 'somewhat unimportant' or 'not important at all'.[32] I create an indicator variable equal to 1 if individuals declare this attribute as being very important and 0 otherwise. Overall, respectively 44%, 60%, 77%, 72%, 57%, 47%, 52%, 61% and 58% of individuals declare the attribute as being 'very important' (see Figure A.17). Figure A.18 shows the share of individuals who declare each attribute as 'very important' as a function of the group they belong to. Each color represents a specific job attribute. 'Industry-specialists' care more about job's opportunities for advancement, job's responsibility and job's salary than other groups. 'Stars' and 'Academia-specialists' share their particular care for job's intellectual challenge and job's degree of independence compared to the other groups. However, 'stars' care significantly less than 'Academia-specialists' about job's security and job's contribution to society. Figure A.19 plots the same figure accounting for PhD major, gender, race and citizenship. Results remain directionally similar.

I now estimate compensating differentials, i.e., the wage premium (or discount) required to make individuals indifferent between Industry and Academia. I model the utility U_i^j of individual i in sector j as:[33]

$$U_i^j = \beta_i \ln(W_i^j) + \alpha_i 1\{j = \text{Industry}\} \tag{3}$$

where α_i captures individual i's preference for Industry relative to Academia and β_i reflects the sensitivity to earnings.[34] Individual i chooses to work in Industry ($j = I$) iff:

$$j(i) = I \iff U_i^I - U_i^A > 0 \iff \beta_i \left(\ln(W_i^I) - \ln(W_i^A)\right) + \alpha_i > 0 \tag{4}$$

Rearranging the expression to solve for the earnings premium that makes the individual indifferent between Industry and Academia:

$$U_i^I - U_i^A = 0 \iff \ln(W_i^I) - \ln(W_i^A) = \frac{-\alpha_i}{\beta_i} \iff \ln(\frac{W_i^I}{W_i^A}) = \frac{-\alpha_i}{\beta_i} \tag{5}$$

where $-\alpha_i/\beta_i$ represents the log earnings premium (or discount) required to compensate individual i for their preference for Academia over Industry. The corresponding earnings premium in dollars, P_i, is then:

$$\ln(\frac{W_i^A + P_i}{W_i^A}) = \frac{-\alpha_i}{\beta_i} \iff \frac{W_i^A + P_i}{W_i^A} = \exp(\frac{-\alpha_i}{\beta_i}) \iff P_i = W_i^A \times (\exp(\frac{-\alpha_i}{\beta_i}) - 1) \tag{6}$$

[32] This question is asked in each SDR edition, so that I have several observations per individual. Note that these variables are imperfect proxy of preferences because these questions are asked to individuals during their career and so might have been influenced by time spent working.

[33] This static model abstracts from explicit modeling of cross-sector career transitions. However, because predicted earnings are estimated over the full observed career, they implicitly incorporate transitions between Academia and Industry. A dynamic extension that models sectoral mobility more explicitly is presented in Appendix E.

[34] When $\alpha_i > 0$, the individual likes Industry more than Academia. When $\alpha_i < 0$, the individual prefers Academia.

This expression shows that the earnings premium required to attract a given individual into Industry depends on three factors: the strength of their non-pecuniary preferences (α_i), their sensitivity to monetary compensation (β_i), and their potential earnings in Academia (W_i^A). The more negative α_i (i.e., the stronger the preference for Academia), the higher the earnings premium required.

Empirically, I have estimated W_i^j for each individual i and sector j. However, because each individual makes only one sector choice, I cannot identify their personal preference parameter α_i. To proceed, I assume that preferences take the form $\alpha_i 1\{j = \text{Industry}\} = \alpha \cdot 1\{j = \text{Industry}\} + \mu_i^j$, where α captures the average preference (or distaste) for Industry relative to Academia, and μ_i^j are i.i.d Type I Extreme Value distributed.

This leads to a discrete choice logit model of the form:

$$Took_i^j = \gamma_0 + \gamma_1 \ln(W_i^j) + \gamma_2 \cdot 1\{j = \text{Industry}\} + \varepsilon_i^j \tag{7}$$

where $Took_i^j$ equal 1 if individual i is observed in sector j, $1\{j = \text{Industry}\}$ equals 1 if the choice relates to Industry and 0 if it relates to Academia. $-\hat{\gamma}_2/\hat{\gamma}_1$ is the empirical estimate of $-\alpha/\beta$ and represents the log earnings premium required to compensate individuals for the average preference for Academia over Industry. To estimate this regression, I transform my dataset to have two rows per individual, one for Industry choice and one for Academia choice. For each individual, I use their (averaged) residualize potential earnings in each sector to proxy for W_i^j.[35]

Table 5 shows the results. Column (1) reports the coefficients from the logit regression. Column (2) reports the marginal effects. The estimated $-\hat{\gamma}_2/\hat{\gamma}_1$ is 0.9. This implies that the average earnings premium is $P = 1.5 \times W_i^A$, meaning that firms would need to pay on average 2.5 times the academic wage to compensate individuals for academic preferences.

To bring further nuance into heterogeneity, I examine how the required earnings premium varies continuously across the propensity score distribution, using a flexible, nonparametric specification where I interact the main variables with deciles of the propensity score. Figure 7 reveals a steeply decreasing premium curve: individuals with a propensity score below 0.1 (40% of the sample) require an earnings premium of approximately 3.8 times the academic wage to be indifferent between sectors, while those with a propensity score above 0.6 (7.6% of the sample) require no premium - and in some cases, appear willing to accept lower earnings to work in Industry. These findings reinforce the role of preferences in shaping sorting patterns and show that the cost of attracting scientists into Industry varies sharply among talent.

[35]For this regression, I use all individuals' earnings, including earnings observations associated with a different sector from the one they initially joined. This approach reflects that I am not solely estimating productivity in each sector separately, but rather examining individuals' career choices, which also accounts for the probability of switching sectors during the career.

I also examine how preferences vary as a function of individuals' estimated earnings potential in Industry - which serves as my proxy for industrial productivity. This lens is particularly important for firms, as it allows to understand what it would take to attract individuals who are most valuable in the industrial setting. It is also helpful in understanding how correlated are preferences with industrial productivity. Figure A.20 plots the estimated earnings premium needed to make individuals indifferent between Industry and Academia against deciles of predicted industrial earnings. The required premium is roughly constant at around 1.5 times the academic wage for individuals in the bottom four deciles of industry earnings potential, but then declines linearly and reaches approximately 0.2 for those in the top decile.[36] This result suggests that preferences are not orthogonal to skills: individuals with high industrial potential also tend to have stronger preferences for Industry. As a consequence, attracting high-productivity scientists is less costly than average estimates of the academic premium would suggest: individuals with above-median industrial productivity require an earnings premium between 1.1 and 0.2 times the academic wage.

7 Robustness

Postdocs - In my analysis, I associate individuals who start a postdoc after PhD graduation as joining Academia. I also reiterate my analysis by associating each individual with a postdoc with the first sector I observe them working in 10 years or more after PhD graduation.[37] The F-statistic for the first-stage equals 25. Using this sector assignment, 71% of postdocs are still observed in Academia and 29% are in Industry. I find a correlation of 0.35 between academic and industrial productivity using potential earnings and a correlation of 0.17 using industrial potential earnings and academic publication potential, similar to the baseline results. 80% of firms' hires have above-median productivity in Industry, similar to the 79% of the baseline results. Universities' hires are similar to the baseline results (see Figure A.21).

Government sector - In my main specification, I bundle individuals who join Academia and Government together. I also check that my results are robust to excluding individuals who join the government sector (2,930 individuals). The F-statistic for the first-stage remains similar (value of 38). I find a similar correlation of 0.33. Figure A.22 shows firms' and universities' hiring composition. 69% of firms' hires have above-median productivity in Industry. Universities' hiring composition is almost unchanged, with 30% of 'stars', 13% of 'Industry-specialists', 23% of 'Academia-specialists' and 34% of 'modest achievers'.

Functional form assumptions - My baseline Marginal Treatment Effects results use the separate method

[36] I also calculate the earnings premium by group: *Stars* have an earnings multiplier of 0.9, *Industry-specialists* of 0.4, *modest achievers* of 1.8 and *Academia-specialist* of 3.7.

[37] 4,357 postdocs (31%) are not observed after 10 years of experience and are therefore dropped from the estimation.

with a polynomial of degree 1 and a probit for the first-stage equation. Figure A.23 shows that the MTE curves are robust to using a polynomial of degree 2 or a semi-parametric approach, which imposes less functional form assumptions. I also check in Figure A.24 that the curves are robust to using a logit or a linear probability model when estimating the first-stage equation.

Empirical strategy - My instrument leverages variation in individuals' graduation year within a field of study. One concern could be changes in skills over time that would correlate with the instrument, e.g. if an increase in firms' demand in Computer Science over time correlates with more Industry-skilled students sorting into that field of study. To test the robustness of my results, I create indicators for each window of 5 years tied to year of graduation (e.g., one indicator for graduation years 1970 to 1974, one indicator for graduation years 1975 to 1979 etc.). I then interact these windows with PhD field fixed-effects and add them into my regressions. This allows me to compare individuals in the same field who graduated around the same time (e.g., Computer Science students between 2000 and 2004) and effectively control for structural major-related shifts in skills composition. The F-statistic decreases to 14 but remains strong. I find a similar correlation of 0.33 between potential earnings in Industry and potential earnings in Academia. Compared to the baseline results, I find that firms only hire 53% of individuals with above-median productivity in Industry (vs 79%), distributed equally between 'stars' and 'Industry-specialists'.

8 Conclusion

Knowledge workers play a central role in driving firm innovation and performance. While individuals' productivity in an industrial setting is unknown at the time of hire, firms can observe measures of their academic productivity. However, because Academia and Industry operate under different institutional logics, it is unclear whether the most academically productive individuals would also be the most productive in Industry. This paper investigates how academic and industrial productivity align and what consequences this alignment - or lack thereof - has for firms' evaluation and attraction of scientific talent. To do so, I use confidential data on PhD graduates and estimate individual-level productivity in both Academia and Industry using Marginal Treatment Effects (MTE), leveraging exogenous variation in private-sector demand within scientific fields over time.

Three key findings emerge. First, academic and industrial productivity are only moderately aligned, with a correlation between 0.19 and 0.28 depending on the productivity measures used. This correlation has decreased over time, implying that academic productivity has become less informative about industrial productivity. Second, I find that 21% of firms' hires have below-median industrial productivity and that these individuals are significantly more likely to be laid off. To examine whether such hires stem from the limited alignment between academic and industrial productivity, I exploit variation in this correlation across majors. I find that low-productivity hires and layoffs are both more

common, and more likely to exhibit strong measures of academic productivity at graduation, in majors where academic and industrial productivity are less aligned. This suggests that weak alignment between academic and industrial productivity creates evaluation constraints, resulting in hires that carry a higher risk of underperformance and layoff. Third, attracting scientists requires substantial variation in compensation: 40% of individuals would require an earnings premium equivalent to 3.8 times the academic wage, while 8% require no premium. Individuals with high industrial potential tend to have weaker preferences for Academia, such that individuals with above-median industrial productivity require an earnings premium between 1.1 and 0.2 times the academic wage.

My results underscore the challenges firms face when hiring productive workers from a different institutional context (Boudou and McKeon, 2024), and provide a theoretical and empirical approach to analyze hiring under this type of cross-context uncertainty. I show that when available measures of productivity reflect a different set of norms and incentives, they can constrain firms' ability to identify individuals who would be productive in their own setting. Importantly, this challenge is not about noisy observable metrics, but reflects a deeper tension in using information generated under one institutional logic to predict performance in another. This tension arises broadly in settings where organizations must assess individuals or ideas across institutional boundaries (Bikard, 2018; Bikard and Marx, 2020; Boudou and Roche, 2023; Cooper et al., 1994). Methodologically, I apply Marginal Treatment Effects to estimate each individual's productivity in both Academia and Industry, decoupled from observed outcomes. I also extend prior work on scientists' compensating differentials by providing a more comprehensive and managerially relevant estimate of the earnings premium required to shift individuals into Industry. While prior work has found that individuals accept on average a 20–25% salary discount to be able to publish (Roach and Sauermann, 2010; Stern, 2004), my approach captures the full bundle of job attributes and accounts for heterogeneity across individuals. More broadly, this paper highlights that leveraging scientific talent is not simply about acquiring PhDs or absorbing external knowledge. Rather, it depends critically on selecting and attracting the right scientists - those whose skills align with the firm's needs (Arora and Gambardella, 1994; Tranchero, 2023). These findings highlight an underexplored dimension of absorptive capacity: the ability to recognize and recruit the individuals best positioned to generate value.

These findings also have significant implications for firms' hiring strategies, especially as organizations increasingly rely on knowledge and innovation to drive growth and performance (Gambardella, 1992). When measures of performance originate from a different institutional context, even highly structured and rational hiring processes can lead to suboptimal outcomes - including costly layoffs and missed opportunities to recruit high-potential talent. The observed variation in alignment between academic and industrial productivity across fields suggests that firms may benefit from field-specific hiring strategies and more nuanced approaches to evaluating academic credentials. Firms may consider supplementing academic indicators with alternative forms of evaluation, such as internships,

collaborative research projects, or structured interviews that simulate real work constraints (Baruffaldi and Poege, 2025; Honoré and Ganco, 2023). As one executive at a major technology company I interviewed said: *"We don't need brilliant thinkers - we need scientists who can move fast, handle ambiguity, and deliver results."*

Understanding the alignment between academic and industrial productivity is critical not just for evaluating candidates, but also for deciding whom to target and how to design effective offers. When firms rely on academic credentials without recognizing how they map onto industrial productivity, they risk overpaying for individuals whose academic strengths do not translate into high productivity in their setting. My results show that inducing high-potential candidates to self-select into Industry is not an easy task, as scientists with high industrial potential vary widely in their preferences and in the compensating differentials they require to enter Industry (Roach and Sauermann, 2010; Stern, 2004). Some may prioritize autonomy or intellectual challenge over salary, while others are more responsive to opportunities for advancement and impact. Firms aiming to attract top scientific talent might therefore benefit from crafting roles and environments that appeal to the specific preferences of their target candidates (Sauermann and Roach, 2014). More broadly, these results suggest that firms cannot rely on one-size-fits-all approaches to hiring scientific talent. Recognizing this complexity and tailoring hiring practices accordingly is likely to be critical for building and sustaining a competitive edge in science-driven industries.

A few limitations of this study provide fertile ground for future work. First, given constraints on the data used, I do not have information about organization-level outcomes, which limits the ability to concretely link scientists' sorting to firm performance. In particular, the study does not address whether the current allocation of talent across sectors is optimal, either from the perspective of firms or from that of social planners. Moreover, I lack information about the characteristics of individuals' publication output, such as journals, references, and abstracts. This limitation restricts my ability to discuss the direction of the innovation that individuals would produce within each sector. Assessing these directional aspects of innovation and how they relate to the types of talent they recruit might be crucial for organizations, especially in firms where scientists have autonomy over the direction of their research. These gaps highlight promising avenues for future research.

References

Agarwal, R. and Ohyama, A. (2013), 'Industry or academia, basic or applied? Career choices and earnings trajectories of scientists', *Management Science* **59**(4), 950–970.

Agrawal, A. K. (2001), 'University-to-industry knowledge transfer: Literature review and unanswered questions', *International Journal of management reviews* **3**(4), 285–302.

Agrawal, A., McHale, J. and Oettl, A. (2017), 'How stars matter: Recruiting and peer effects in evolutionary biology', *Research Policy* **46**(4), 853–867.

Andresen, M. E. (2018), 'Exploring marginal treatment effects: Flexible estimation using stata', *The Stata Journal* **18**(1), 118–158.

Arellano-Bover, J. (2024), 'Career consequences of firm heterogeneity for young workers: First job and firm size', *Journal of Labor Economics* **42**(2), 549–589.

Arora, A., Belenzon, S., Kosenko, K., Suh, J. and Yafeh, Y. (2024), 'The rise of scientific research in corporate America', *Organization Science* .

Arora, A., Belenzon, S. and Patacconi, A. (2018), 'The decline of science in corporate R&D', *Strategic Management Journal* **39**(1), 3–32.

Arora, A., Belenzon, S., Suh, J. and Zhang, H. (2025), 'The rise of absorptive research in corporate America: 1945-1980', *NBER Working Paper* .

Arora, A. and Gambardella, A. (1994), 'Evaluating technological information and utilizing it: Scientific knowledge, technological capability, and external linkages in biotechnology', *Journal of Economic Behavior & Organization* **24**(1), 91–114.

Azoulay, P., Ganguli, I. and Zivin, J. G. (2017), 'The mobility of elite life scientists: Professional and personal determinants', *Research Policy* **46**(3), 573–590.

Azoulay, P., Graff Zivin, J. S. and Manso, G. (2011), 'Incentives and creativity: Evidence from the academic life sciences', *The RAND Journal of Economics* **42**(3), 527–554.

Barney, J. (1991), 'Firm resources and sustained competitive advantage', *Journal of Management* **17**(1), 99–120.

Baruffaldi, S. and Poege, F. (2025), 'Like stars: How firms learn at scientific conferences', *Management Science* **71**(3), 2056–2078.

Becker, G. S. (1962), 'Investment in human capital: A theoretical analysis', *Journal of Political Economy* **70**(5, Part 2), 9–49.

Bidwell, M. (2011), 'Paying more to get less: The effects of external hiring versus internal mobility', *Administrative Science Quarterly* **56**(3), 369–407.

Bidwell, M. and Mollick, E. (2015), 'Shifts and ladders: Comparing the role of internal and external mobility in managerial careers', *Organization Science* **26**(6), 1629–1645.

Bikard, M. (2018), 'Made in academia: The effect of institutional origin on inventors' attention to science', *Organization Science* **29**(5), 818–836.

Bikard, M. and Marx, M. (2020), 'Bridging academia and industry: How geographic hubs connect university science and corporate technology', *Management Science* **66**(8), 3425–3443.

Bikard, M., Vakili, K. and Teodoridis, F. (2019), 'When collaboration bridges institutions: The impact of university–industry collaboration on academic productivity', *Organization Science* **30**(2), 426–445.

Björklund, A. and Moffitt, R. (1987), 'The estimation of wage gains and welfare gains in self-selection models', *The Review of Economics and Statistics* pp. 42–49.

Boudou, J. and McKeon, J. (2024), 'Innovation under resource constraints: Supercomputing in scientific research', *Working Paper* .

Boudou, J. and Roche, M. P. (2023), 'Bringing science to market: Knowledge foundations and performance', *Working Paper* .

Brinch, C. N., Mogstad, M. and Wiswall, M. (2017), 'Beyond late with a discrete instrument', *Journal of Political Economy* **125**(4), 985–1039.

Campbell, B. A., Coff, R. and Kryscynski, D. (2012), 'Rethinking sustained competitive advantage from human capital', *Academy of Management Review* **37**(3), 376–395.

Campbell, B. A., Ganco, M., Franco, A. M. and Agarwal, R. (2012), 'Who leaves, where to, and why worry? Employee mobility, entrepreneurship and effects on source firm performance', *Strategic Management Journal* **33**(1), 65–87.

Cappelli, P. (2019), 'Your approach to hiring is all wrong', *Harvard Business Review* **97**(3), 48–58.

Carnahan, S., Agarwal, R. and Campbell, B. A. (2012), 'Heterogeneity in turnover: The effect of relative compensation dispersion of firms on the mobility and entrepreneurship of extreme performers', *Strategic Management Journal* **33**(12), 1411–1430.

Carneiro, P., Heckman, J. J. and Vytlacil, E. J. (2011), 'Estimating marginal returns to education', *American Economic Review* **101**(6), 2754–2781.

Cockburn, I. M. and Henderson, R. M. (1998), 'Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery', *The Journal of Industrial Economics* **46**(2), 157–182.

Coff, R. and Kryscynski, D. (2011), 'Drilling for micro-foundations of human capital-based competitive advantages. invited editorial', *Journal of Management* **37**(5), 1429–1443.

Coff, R. and Rickley, M. (2021), 'Strategic human capital: Fit for the future', *Strategic Management: State of the Field and its Future; Duhaime, IM, Hitt, MA, Lyles, MA, Eds* pp. 579–593.

Coff, R. W. (1997), 'Human assets and management dilemmas: Coping with hazards on the road to resource-based theory', *Academy of Management Review* **22**(2), 374–402.

Cohen, W. M. and Levinthal, D. A. (1990), 'Absorptive capacity: A new perspective on learning and innovation', *Administrative Science Quarterly* **35**(1), 128–152.

Conti, A. and Visentin, F. (2015), 'A revealed preference analysis of PhD students' choices over employment outcomes', *Research Policy* **44**(10), 1931–1947.

Cooper, A. C., Gimeno-Gascon, F. J. and Woo, C. Y. (1994), 'Initial human and financial capital as predictors of new venture performance', *Journal of Business Venturing* **9**(5), 371–395.

Cornelissen, T., Dustmann, C., Raute, A. and Schönberg, U. (2016), 'From LATE to MTE: Alternative methods for the evaluation of policy interventions', *Labour Economics* **41**, 47–60.

Dasgupta, P. and David, P. A. (1994), 'Toward a new economics of science', *Research Policy* **23**(5), 487–521.

Deming, D. J. and Noray, K. L. (2018), *STEM careers and technological change*, Vol. 24, National Bureau of Economic Research Cambridge, MA.

Ding, W. W., Levin, S. G., Stephan, P. E. and Winkler, A. E. (2010), 'The impact of information technology on academic scientists' productivity and collaboration patterns', *Management Science* **56**(9), 1439–1461.

Dobbie, W., Goldin, J. and Yang, C. S. (2018), 'The effects of pre-trial detention on conviction, future crime, and employment: Evidence from randomly assigned judges', *American Economic Review* **108**(2), 201–240.

Gambardella, A. (1992), 'Competitive advantages from in-house scientific research: The US pharmaceutical industry in the 1980s', *Research Policy* **21**(5), 391–407.

Gambardella, A., Panico, C. and Valentini, G. (2015), 'Strategic incentives to human capital', *Strategic Management Journal* **36**(1), 37–52.

Ganco, M. (2013), 'Cutting the gordian knot: The effect of knowledge complexity on employee mobility and entrepreneurship', *Strategic Management Journal* **34**(6), 666–686.

Ganco, M., Ziedonis, R. H. and Agarwal, R. (2015), 'More stars stay, but the brightest ones still leave: Job hopping in the shadow of patent enforcement', *Strategic Management Journal* **36**(5), 659–685.

Gittelman, M. and Kogut, B. (2003), 'Does good science lead to valuable knowledge? Biotechnology firms and the evolutionary logic of citation patterns', *Management Science* **49**(4), 366–382.

Goolsbee, A. and Syverson, C. (2023), 'Monopsony power in higher education: A tale of two tracks', *Journal of Labor Economics* **41**(S1), S257–S290.

Grennan, M., Myers, K., Swanson, A. and Chatterji, A. (2024), 'No free lunch? Welfare analysis of firms selling through expert intermediaries', *Review of Economic Studies* p. rdae090.

Heckman, J. J. (1981), 'An incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process', *The Structural Analysis of Discrete Data* pp. 179–195.

Heckman, J. J. and Vytlacil, E. J. (1999), 'Local instrumental variables and latent variable models for identifying and bounding treatment effects', *Proceedings of the National Academy of Sciences* **96**(8), 4730–4734.

Heckman, J. J. and Vytlacil, E. J. (2007), 'Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments', *Handbook of Econometrics* **6**, 4875–5143.

Hess, A. M. and Rothaermel, F. T. (2011), 'When are assets complementary? Star scientists, strategic alliances, and innovation in the pharmaceutical industry', *Strategic Management Journal* **32**(8), 895–909.

Hill, R. and Stein, C. (2025), 'Race to the bottom: Competition and quality in science', *The Quarterly Journal of Economics* p. qjaf010.

Hilmer, M. J., Ransom, M. R. and Hilmer, C. E. (2015), 'Fame and the fortune of academic economists: How the market rewards influential research in economics', *Southern Economic Journal* **82**(2), 430–452.

Hiraiwa, T., Braguinsky, S. and Agarwal, R. (2025), 'The evolution of talent allocation into Academia: Institution-building and graduates' choices during Japan's industrialization', *Working Paper* .

Honoré, F. and Ganco, M. (2023), 'Entrepreneurial teams' acquisition of talent: Evidence from technology manufacturing industries using a two-sided approach', *Strategic Management Journal* **44**(1), 141–170.

Hsu, D. H. and Kuhn, J. M. (2023), 'Academic stars and licensing experience in university technology commercialization', *Strategic Management Journal* **44**(3), 887–905.

Huang, F. and Cappelli, P. (2010), 'Applicant screening and performance-related outcomes', *American Economic Review* **100**(2), 214–218.

Kaiser, U., Kongsted, H. C., Laursen, K. and Ejsing, A.-K. (2018), 'Experience matters: The role of academic scientist mobility for industrial innovation', *Strategic Management Journal* **39**(7), 1935–1958.

Katz, D. A. (1973), 'Faculty salaries, promotions, and productivity at a large university', *American Economic Review* **63**(3), 469–477.

Kryscynski, D. (2021), 'Firm-specific worker incentives, employee retention, and wage–tenure slopes', *Organization Science* **32**(2), 352–375.

Lazear, E. P. (2009), 'Firm-specific human capital: A skill-weights approach', *Journal of Political Economy* **117**(5), 914–940.

Levin, S. G. and Stephan, P. E. (1991), 'Research productivity over the life cycle: Evidence for academic scientists', *American Economic Review* pp. 114–132.

Liu, H., Mihm, J. and Sosa, M. E. (2018), 'Where do stars come from? The role of star vs. nonstar collaborators in creative settings', *Organization Science* **29**(6), 1149–1169.

Marx, M., Strumsky, D. and Fleming, L. (2009), 'Mobility, skills, and the michigan non-compete experiment', *Management Science* **55**(6), 875–889.

Merton, R. K. (1973), *The sociology of science: Theoretical and empirical investigations*, University of Chicago press.

Ng, W. and Stuart, T. E. (2022), 'Acquired employees versus hired employees: Retained or turned over?', *Strategic Management Journal* **43**(5), 1025–1045.

Oettl, A. (2012), 'Reconceptualizing stars: Scientist helpfulness and peer performance', *Management Science* **58**(6), 1122–1140.

Oyer, P. (2006), 'Initial labor market conditions and long-term outcomes for economists', *Journal of Economic Perspectives* **20**(3), 143–160.

Palomeras, N. and Melero, E. (2010), 'Markets for inventors: Learning-by-hiring as a driver of mobility', *Management Science* **56**(5), 881–895.

Perkmann, M., Tartari, V., McKelvey, M., Autio, E., Broström, A., D'este, P., Fini, R., Geuna, A., Grimaldi, R., Hughes, A. et al. (2013), 'Academic engagement and commercialisation: A review of the literature on university–industry relations', *Research Policy* **42**(2), 423–442.

Polidoro Jr, F. and Theeke, M. (2012), 'Getting competition down to a science: The effects of technological competition on firms' scientific publications', *Organization Science* **23**(4), 1135–1153.

Prendergast, C. (1999), 'The provision of incentives in firms', *Journal of Economic Literature* **37**(1), 7–63.

Roach, M. and Sauermann, H. (2010), 'A taste for science? PhD scientists' academic orientation and self-selection into research careers in industry', *Research policy* **39**(3), 422–434.

Roach, M. and Sauermann, H. (2024), 'Can technology startups hire talented early employees? Ability, preferences, and employee first job choice', *Management Science* **70**(6), 3619–3644.

Rosen, S. (1981), 'The economics of superstars', *American Economic Review* **71**(5), 845–858.

Sampat, B. and Williams, H. L. (2019), 'How do patents affect follow-on innovation? Evidence from the human genome', *American Economic Review* **109**(1), 203–236.

Sauermann, H. and Roach, M. (2014), 'Not all scientists pay to be scientists: Phds' preferences for publishing in industrial employment', *Research Policy* **43**(1), 32–47.

Sauermann, H. and Stephan, P. (2013), 'Conflicting logics? A multidimensional view of industrial and academic science', *Organization Science* **24**(3), 889–909.

Sevcenko, V. and Ethiraj, S. (2018), 'How do firms appropriate value from employees with transferable skills? A study of the appropriation puzzle in actively managed mutual funds', *Organization Science* **29**(5), 775–795.

Shvadron, D. (2023), 'Bread upon the waters: Corporate science and the benefits from follow-on public research', *Working Paper* .

Singh, J. and Agrawal, A. (2011), 'Recruiting for ideas: How firms exploit the prior inventions of new hires', *Management Science* **57**(1), 129–150.

Starr, E., Frake, J. and Agarwal, R. (2019), 'Mobility constraint externalities', *Organization Science* **30**(5), 961–980.

Starr, E., Ganco, M. and Campbell, B. A. (2018), 'Strategic human capital management in the context of cross-industry and within-industry mobility frictions', *Strategic Management Journal* **39**(8), 2226–2254.

Stephan, P. (2012), *How economics shapes science*, Harvard University Press.

Stern, S. (2004), 'Do scientists pay to be scientists?', *Management Science* **50**(6), 835–853.

Stuart, T. E. and Ding, W. W. (2006), 'When do scientists become entrepreneurs? the social structural antecedents of commercial activity in the academic life sciences', *American journal of sociology* **112**(1), 97–144.

Tartari, V. and Breschi, S. (2012), 'Set them free: scientists' evaluations of the benefits and costs of university–industry research collaboration', *Industrial and Corporate Change* **21**(5), 1117–1147.

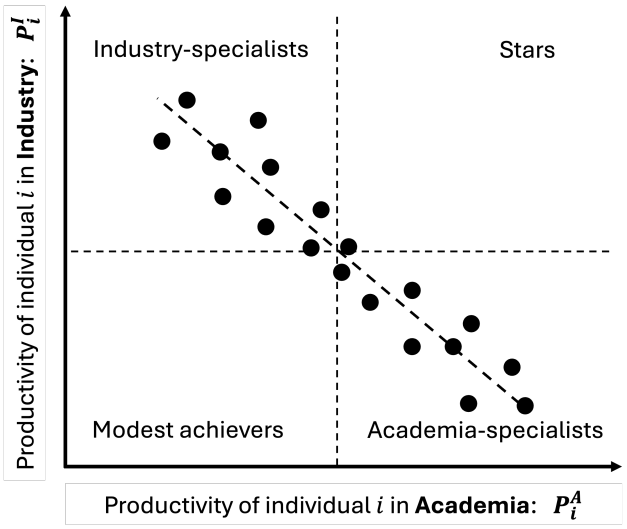
Toole, A. A. and Czarnitzki, D. (2010), 'Commercializing science: Is there a university "brain drain" from academic entrepreneurship?', *Management Science* **56**(9), 1599–1614.

Tranchero, M. (2023), 'Finding diamonds in the rough: Data-driven opportunities and pharmaceutical innovation', *Working Paper* .

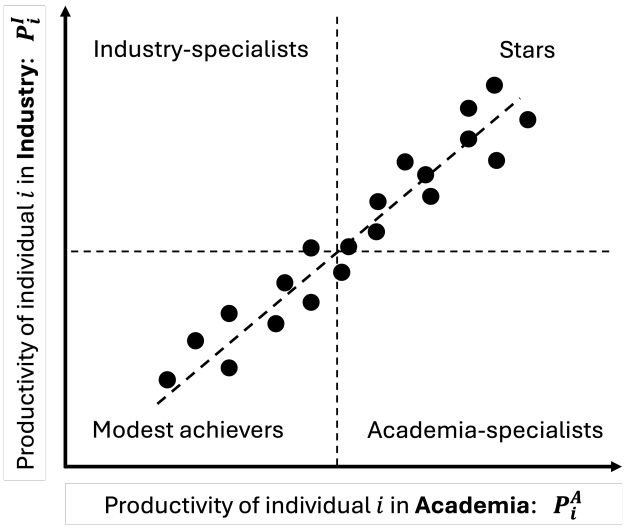
Yue, D. (2024), 'I, Google: Estimating the Impact of Corporate Involvement on AI Research', *Working Paper* .

Figures

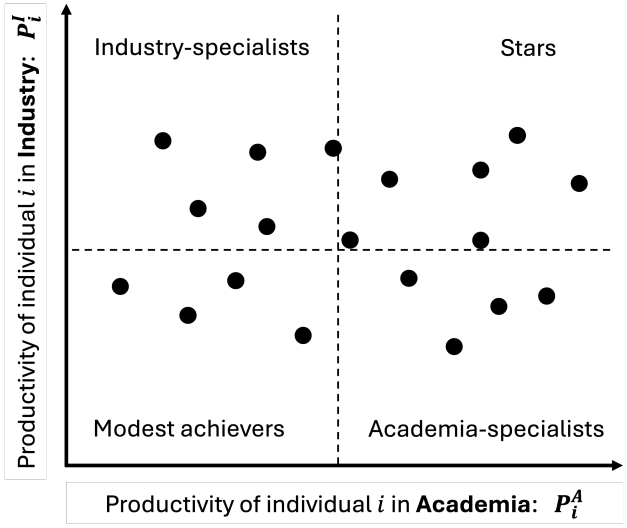
Figure 1: Stylized joint distributions of academic and industrial productivity



(a) Negative correlation



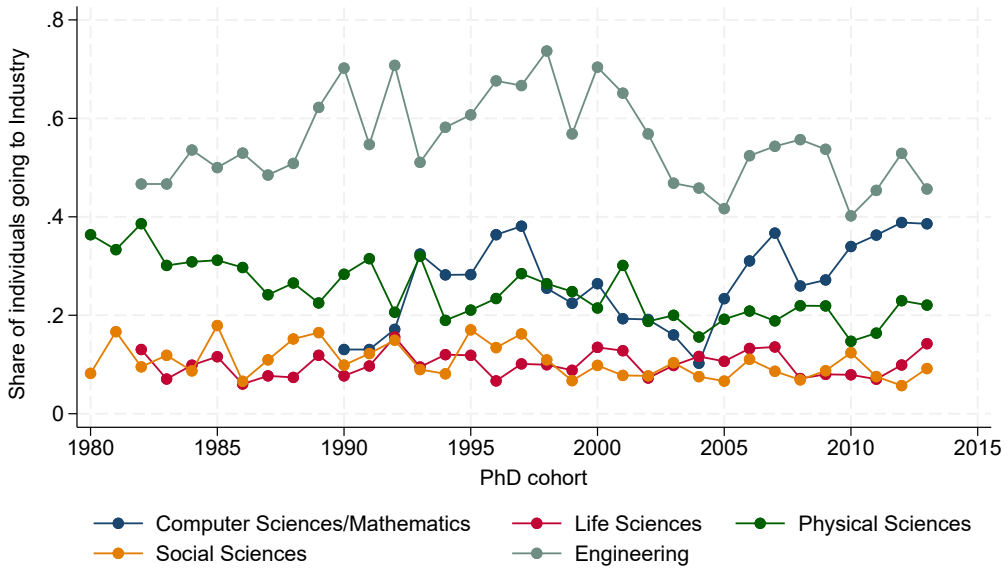
(b) Positive correlation



(c) Null correlation

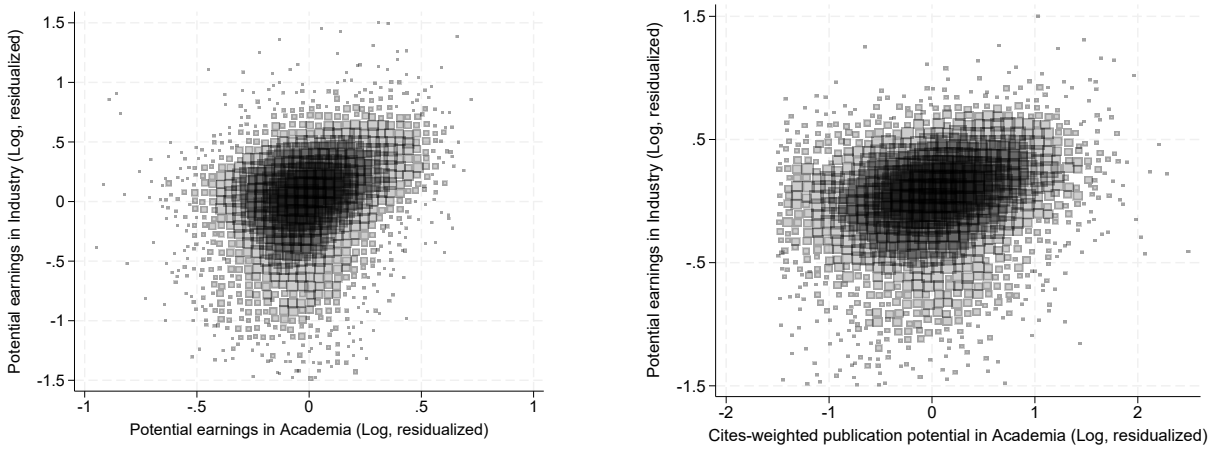
Notes: Each dot represents an individual. Each panel illustrates a stylized joint distribution between academic and industrial productivity. (a) Negative correlation: individuals tend to be productive in one sector but not the other. (b) Positive correlation: individuals productive in one sector are also productive in the other. (c) Zero correlation: no systematic relationship between productivity across sectors.

Figure 2: Sorting patterns, by PhD cohort and field



Notes: This figure shows the share of individuals entering Industry by PhD graduation cohort and PhD field. Dots are omitted for early cohorts in certain fields due to confidentiality restrictions.

Figure 3: Relationship between individuals' industrial and academic productivity

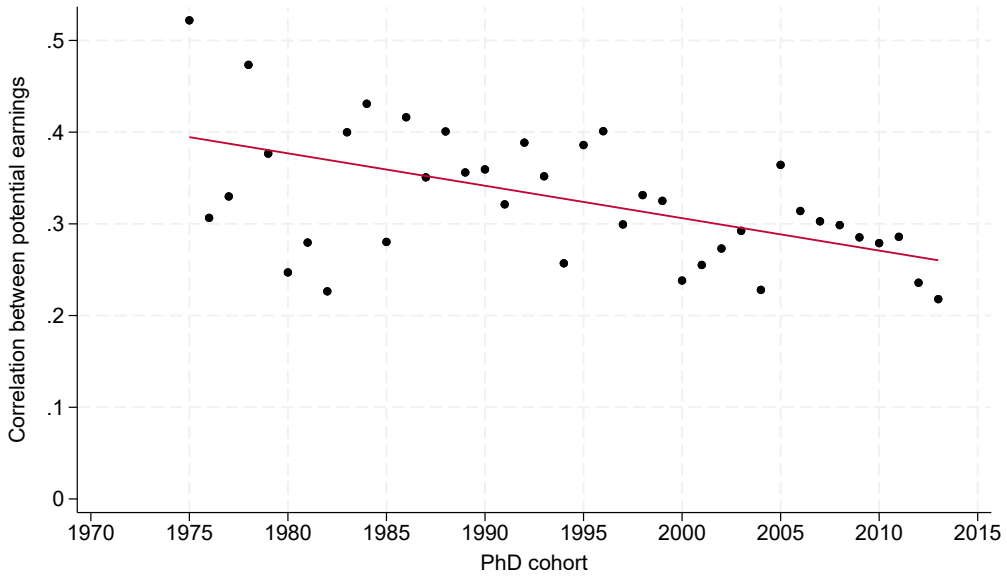


(a) Using predicted earnings

(b) Using predicted earnings and publications

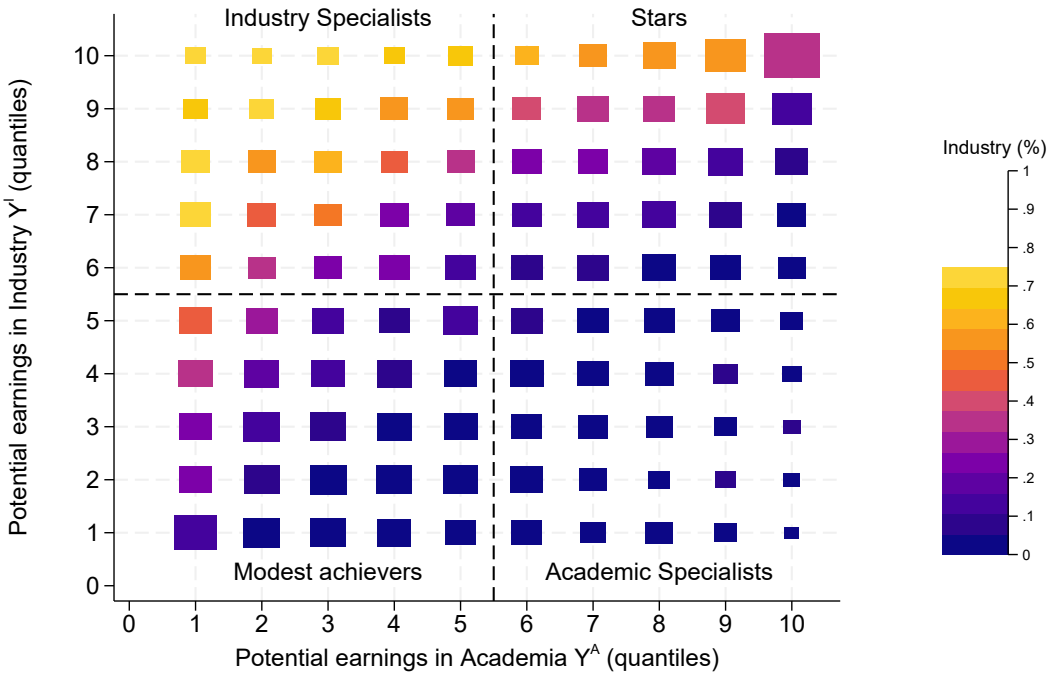
Notes: This figure shows a scatter plot of estimated productivity in Academia (x-axis) and estimated productivity in Industry (y-axis), based on Marginal Treatment Effect (MTE) estimates. Panel (a): Individuals' productivity in Academia and Industry is proxied by their predicted earnings in each sector (in thousands of 2015 USD), residualized by experience and its square, and averaged at the individual level. Panel (b): Individuals' industrial productivity is proxied by their predicted earnings in Industry (in thousands of 2015 USD), and individuals' academic productivity is proxied by their predicted cumulative publication output in Academia, weighted by citations received in the first five years. Predicted earnings are residualized by experience and its square, and averaged at the individual level. Predicted publications are residualized by experience, its square, and PhD major.

Figure 4: Correlation between academic and industrial productivity, by PhD cohort



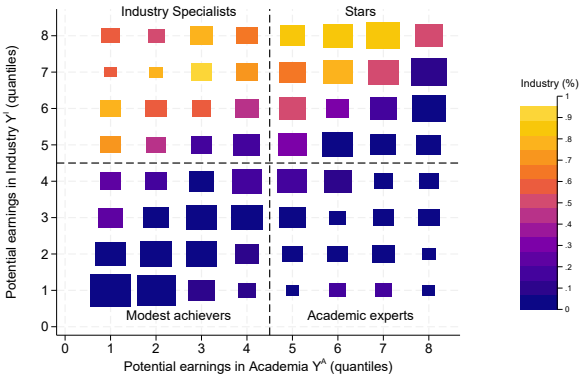
Notes: This figure shows the correlation between individuals' academic and industrial productivity, as a function of PhD cohort of graduation. Individuals' productivity in Academia and Industry is proxied by their predicted earnings in each sector.

Figure 5: Sorting patterns by academic and industrial productivity

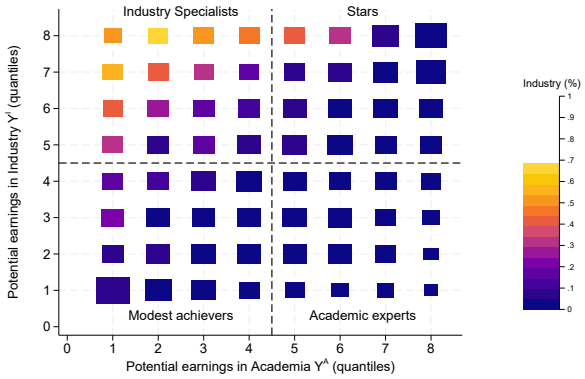


Notes: Each cell in the grid represents a group of individuals based on their decile of academic productivity (x-axis) and industrial productivity (y-axis). Productivity in each sector is measured using predicted potential earnings residualized by experience and its square, and averaged at the individual level. The size of each square is proportional to the number of individuals in that cell. The color intensity indicates the share of individuals in that group who are observed to join Industry (lighter = more Industry, darker = more Academia).

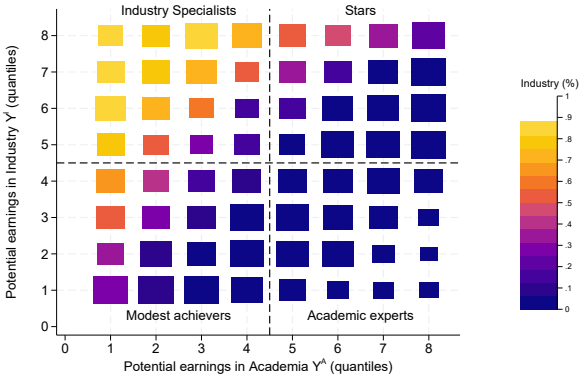
Figure 6: Sorting patterns by academic and industrial productivity, by PhD field



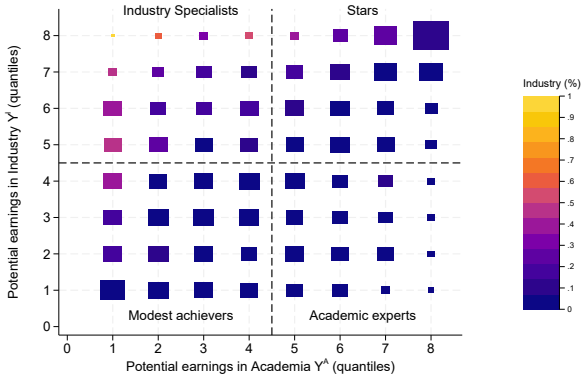
(a) Computer and Mathematical Sciences



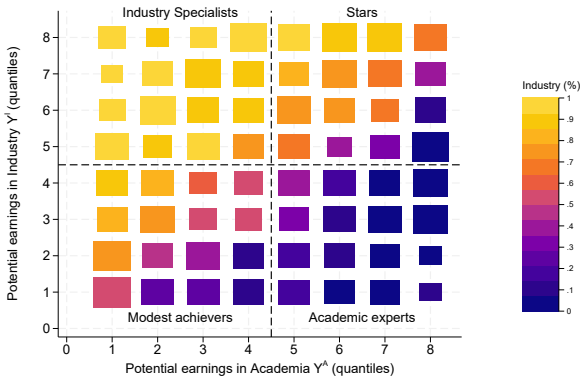
(b) Life Sciences



(c) Physical Sciences



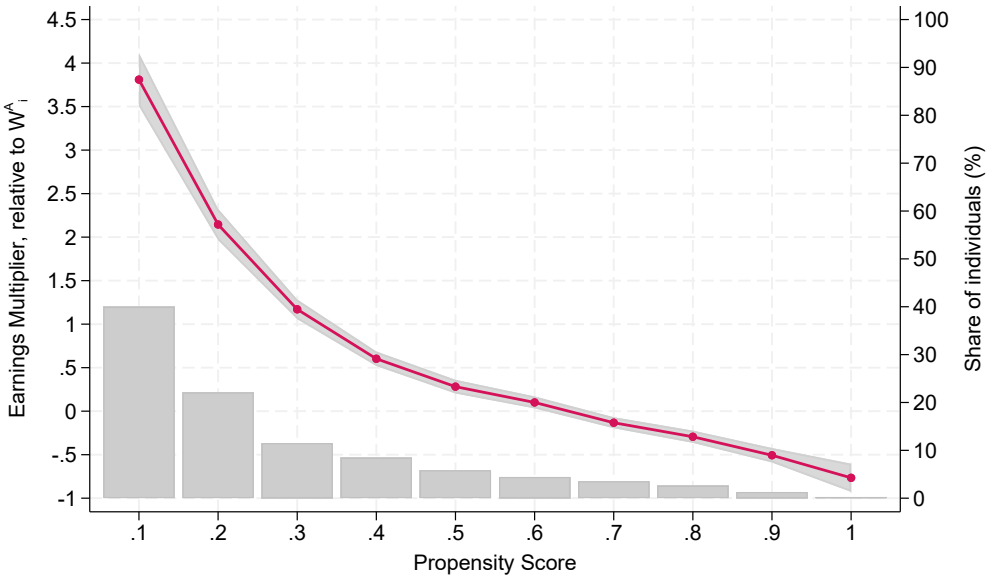
(d) Social Sciences



(e) Engineering

Notes: Each panel shows sorting patterns for a specific PhD field. Within each panel, each cell in the grid represents a group of individuals based on their decile of academic productivity (x-axis) and industrial productivity (y-axis). Productivity is measured using predicted potential earnings, residualized by experience and its square and averaged at the individual level. The size of each square is proportional to the number of individuals in that cell. The color intensity indicates the share of individuals in that group who are observed to join Industry (lighter = more Industry, darker = more Academia).

Figure 7: Estimated compensating differential, by propensity to choose Industry



Notes: This figure shows the estimated earnings premium required to make individuals indifferent between Industry and Academia, by decile of the propensity score. The x-axis divides the propensity score (predicted probability of entering Industry based on observables and the instrument) into ten bins (0–0.1, 0.1–0.2, ..., 0.9–1). For each bin, the line shows the estimated compensating differential relative to academic earnings, calculated as $\exp(\frac{-\alpha_k}{\beta}) - 1$ where α_k captures average preferences for Industry in bin k and β reflects sensitivity to earnings. The shaded area indicates 95% confidence intervals, based on boostrapped standard errors clustered at the individual level. Bar heights represent the share of individuals in each bin of the propensity score distribution.

Tables

Table 1: Summary statistics

	(1) Academia		(2) Industry		(1)-(2) Difference	
	mean	sd	mean	sd	diff in means	t-stat
Earnings (th., 2015 USD)	99.9	58.2	167.4	102.6	-67.48	-43.4***
Publications (cum.)	20.8	30.4	5.3	11.2	15.50	49.3***
PhD Graduation Year	2000.4	10.5	2000.6	9.9	-0.11	-0.7
White	0.7	0.5	0.6	0.5	0.11	13.3***
Female	0.4	0.5	0.3	0.4	0.18	23.6***
Mother has PhD	0.1	0.2	0.1	0.2	0.01	2.8**
Father has PhD	0.2	0.4	0.2	0.4	0.03	5.0***
American	0.8	0.4	0.6	0.5	0.13	16.4***
Observations	17,895		4,714		22,609	

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. This table lists summary statistics for the sample of 22,609 individuals, separately for individuals who join Academia and individuals who join Industry. Column 'diff in means' reports the difference between the mean in Academia and the mean in Industry. Column 't-stat' reports the t-statistic of a test of equality of means. *Earnings*: Average annual earnings in thousands of USD 2015, computed across all survey waves in which the individual is observed. *Publications*: aggregate number of publications between PhD graduation and 2017. *PhD graduation year*: Fiscal year of doctorate. *White*: 0/1=1 if the individual's race/ethnicity is reported as white. *Female*: 0/1=1 if the individual reports being a female. *Mother has PhD*: 0/1=1 if the individual's mother highest educational attainment is a PhD. *Father has PhD*: 0/1=1 if the individual's father highest educational attainment is a PhD. *American*: 0/1=1 if the individual is a U.S. citizen.

Table 2: Tests of validity of the instrument

Test of:	Strategic Graduation			Strategic Entry	
	PhD Duration	$<$ Average	$>$ Average	Z_i	Industry
	(1)	(2)	(3)	(4)	(5)
Z_i	-0.186	-0.030	0.047		
	(0.258)	(0.085)	(0.084)		
Z_i (entry)				-0.030	0.003
				(0.022)	(0.043)
Demographics	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes
PhD major FE	Yes	Yes	Yes	Yes	Yes
Doct. inst. FE	Yes	Yes	Yes	Yes	Yes
Macro cond.	Yes	Yes	Yes	Yes	Yes
Observations	16,733	16,720	16,720	14,279	14,279
R-sq	0.14	0.18	0.18	0.84	0.25

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The outcome in Column (1) is calculated as the difference between PhD graduation year and PhD entry year. The outcome in Column (2) is an indicator equal to 1 if PhD duration is strictly below the average PhD duration of other individuals in the same cohort and PhD major. The outcome in Column (3) is an indicator equal to 1 if PhD duration is strictly above the average PhD duration of other individuals in the same cohort and PhD major. The outcome in Column (4) is the instrument, calculated as in Equation 1. The outcome in Column (5) is an indicator equal to 1 if the individual joins Industry. Z_i *(entry)* is the instrument calculated for the year before PhD entry. *Demographics* includes an indicator equal to 1 if the individual is female, an indicator equal to 1 if the individual is white, an indicator equal to 1 if at least one of the individual's parents has a doctoral degree, an indicator for being an American citizen and birth place fixed-effects. *Experience* includes a linear and quadratic terms for experience, calculated as the difference between survey year and PhD graduation year. *Macro cond.* includes controls for the unemployment rate and field-specific federal funding at the time of graduation. Standard errors are clustered at the PhD major$\times$PhD cohort level.

Table 3: In majors with weaker alignment between academic and industrial productivity, firms hire individuals who are more likely to be laid-off and to have low industrial productivity

	% layoff	% below-median P_i^I
	(1)	(2)
Acad.-Ind. correlation	-0.026***	-0.294***
	(0.003)	(0.005)
Observations	56	56
Mean of dep. var.	0.10	0.15

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The level of analysis is a PhD major. For each PhD major, the outcome in Column (1) is the share of individuals who joined Industry and were ever laid off. For each PhD major, the outcome in Column (2) is the share of individuals who joined Industry whose predicted industrial productivity falls below the median. The key independent variable in both specifications is the correlation between academic and industrial productivity within each PhD major. Lower correlation reflects weaker alignment between the two sectors. All regressions are weighted by the number of individuals in each major. Standard errors are clustered at the PhD major×PhD cohort level.

Table 4: In majors with weaker alignment between academic and industrial productivity, laid-off and low-performing hires are more likely to exhibit indicators of academic productivity at graduation

	z-publication (#)		z-publication (0/1)		Top PhD Institution	
	(1)	(2)	(3)	(4)	(5)	(6)
Acad.-Ind. correlation	-1.164***	-0.778**	-0.765*	-1.050***	-0.084	-0.457***
	(0.411)	(0.324)	(0.444)	(0.337)	(0.183)	(0.093)
Demographics FE	Yes	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes	Yes
Observations	496	712	496	712	714	1,114
Mean of dep. var.	-0.11	-0.11	-0.03	0.007	0.25	0.10

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. In Columns (1), (3), (5), the sample is restricted to individuals who joined Industry and were ever laid off. In Columns (2), (4), (6), the sample is restricted to individuals who joined Industry whose predicted industrial productivity falls below the median. In Columns (1) and (2), the outcome is an indicator for having at least one publication at graduation. This measure is standardized by PhD major. In Columns (3) and (4), the outcome is publication count at graduation. This measure is standardized by PhD major. In Columns (5) and (6), the outcome is an indicator for graduating from a top-20 U.S. PhD program. The key independent variable is the correlation between academic and industrial productivity within each PhD major. Lower correlation reflects weaker alignment between the two sectors. All regressions control for demographic characteristics and PhD cohort fixed-effects. Standard errors are clustered at the PhD major×PhD cohort level.

Table 5: Estimation of average compensating differentials between Academia and Industry

	Choose Sector=1	
	Logit	Marginal Effect
	(1)	(2)
Log(Earnings)	3.386***	0.505***
	(0.053)	(0.008)
Sector (Industry=1)	-3.093***	-0.461***
	(0.036)	(0.001)
Constant	1.438***	
	(0.018)	
Observations	43,632	43,632

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The unit of observation is the individual-sector pair. Each individual appears twice: once for Industry and once for Academia. The dependent variable equals 1 if the individual chose the sector represented by that row. *Log(Earnings)* refers to the log of estimated potential earnings in the corresponding sector. *Sector (Industry=1)* is a dummy equal to 1 if the row corresponds to Industry. Column (2) reports marginal effects. Robust standard errors are clustered at the individual level.

Not All That Glitters is Gold: Firm Hiring in the Market for Knowledge Workers

Appendix

Appendix A Figures

Figure A.1: PhD fields and majors

Field

Major

Computer Sciences/Mathematics
Computer/Information Sciences
Applied Mathematics
Operations Research
Statistics
Other Mathematics

Life Sciences
Biochemistry and Biophysics
Biology
Botany
Cell and Molecular Biology
Ecology
Genetics
Microbiological Sciences
Animal Sciences
Food Sciences
Plant Sciences
Other Agricultural Sciences
Pharmacology
Zoology
Other biological Sciences
Physiology and Pathology
Environmental Sciences
Forestry Sciences

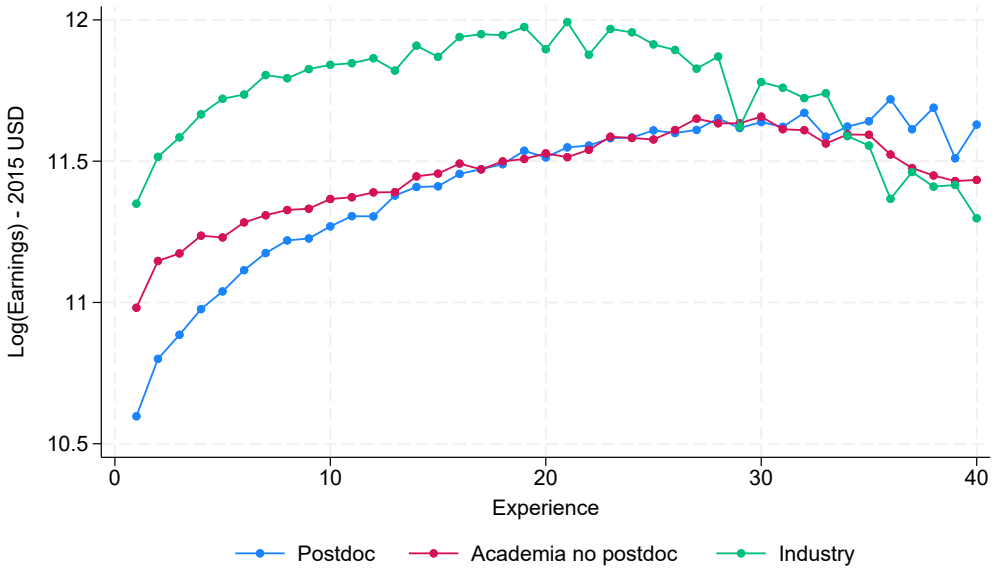
Physical Sciences
Astronomy/Astrophysics
Physics
Other physical sciences
Chemistry
Atmospheric Sciences
Geology
Geological Sciences

Engineering
Aerospace Engineering
Chemical Engineering
Civil Engineering
Computer Engineering
Electrical Engineering
Industrial Engineering
Mechanical Engineering
Bioengineering and biomedical Engineering
Engineering Sciences
Environmental Engineering
Materials Engineering
Other Engineering

Social Sciences
Agricultural economics
Economics
Public Policy
Political Sciences
Educational Psychology
Clinical Psychology
Counseling psychology
Experimental psychology
Industrial/Organizational psychology
Social psychology
Other psychology
Anthropology/Archaeology
Criminology
Sociology
Area and Ethnic Studies
Geography
Other social sciences

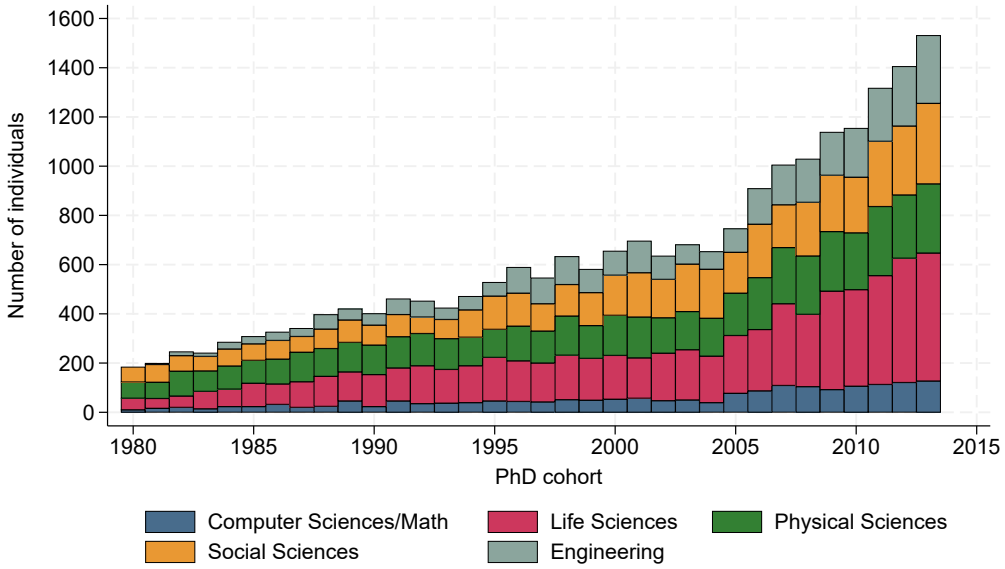
Notes: This figure shows the PhD fields and majors present in my sample. For confidentiality reasons, I am not allowed to export the number of students per PhD major, but Figure A.3 shows the number of individuals by PhD cohort of graduation and main PhD field of study.

Figure A.2: Earnings curve by status after PhD graduation



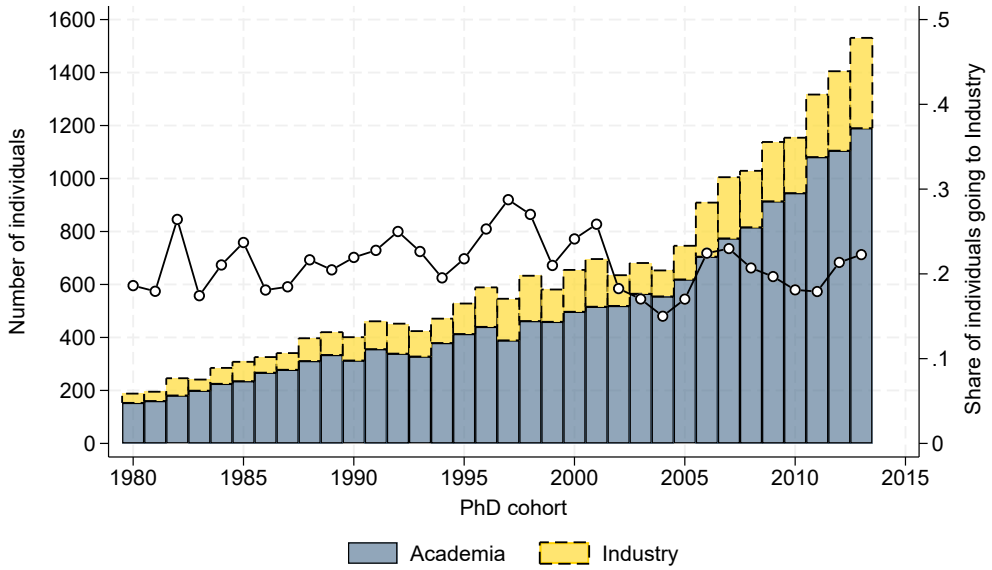
Notes: This figure shows the earnings profile of individuals as a function of their employment status upon PhD graduation. The blue curve represents individuals who start a postdoc after PhD graduation. The red curve represents individuals who join Academia after PhD graduation but not under a postdoc position. The green curve represents individuals who join the private sector. Earnings are expressed in 2015 USD, winsorized below the 1^{st} and above the 99^{th} percentiles and logged.

Figure A.3: Supply, by PhD cohort and field



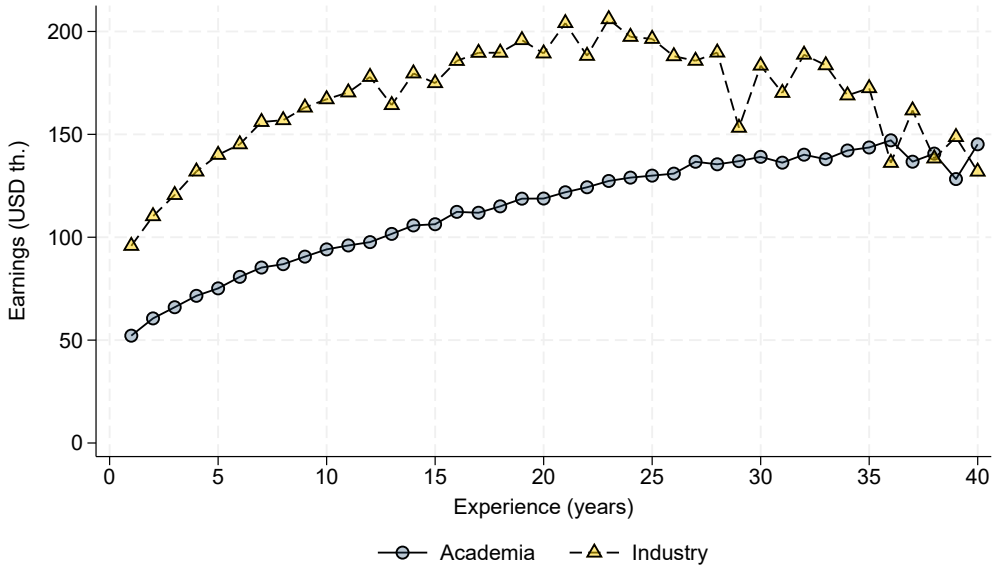
Notes: This figure shows the number of individuals by PhD cohort of graduation and main PhD field of study. The values for years 1980 and 1981 for the field of Engineering are combined into a single dot corresponding to the year 1981 for confidentiality reasons.

Figure A.4: Sorting patterns, by PhD cohort



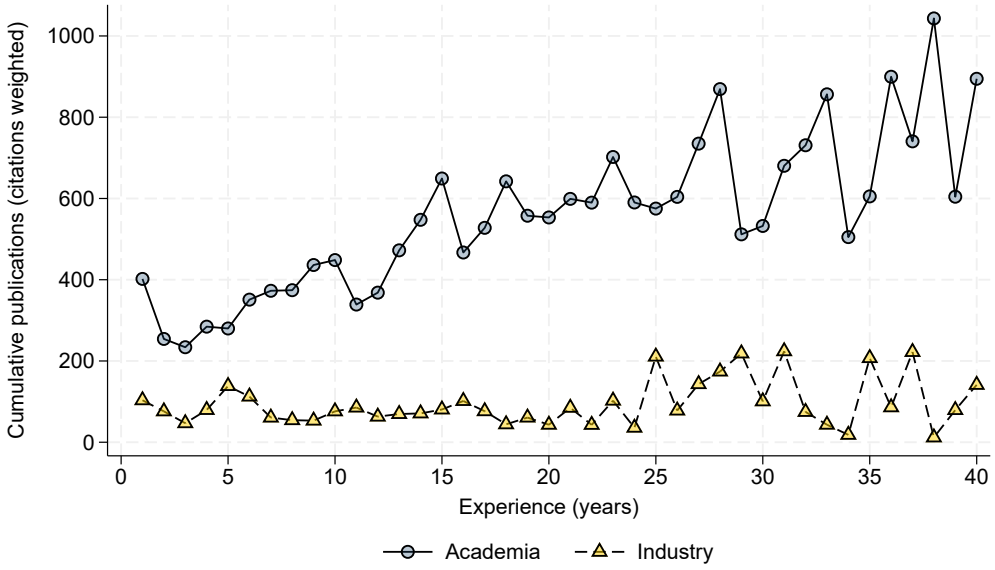
Notes: This figure shows the number of individuals entering Academia (blue bars) and Industry (yellow bars) by PhD graduation cohort. The black line shows the share entering Industry.

Figure A.5: Earnings as a function of experience and sector joined at graduation



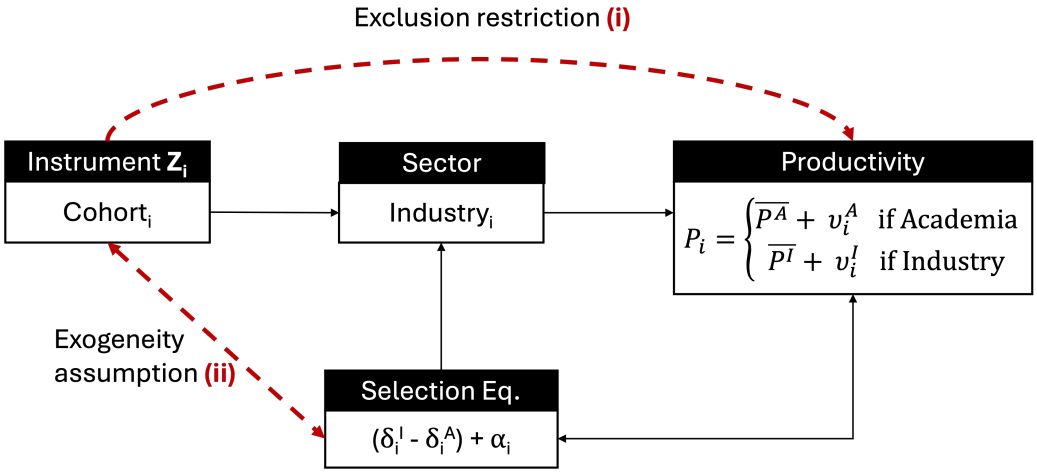
Notes: This figure shows average earnings in thousands of USD 2015 as a function of the sector where individuals started their career.

Figure A.6: Citation-weighted publications (stock) as a function of experience and sector joined at graduation



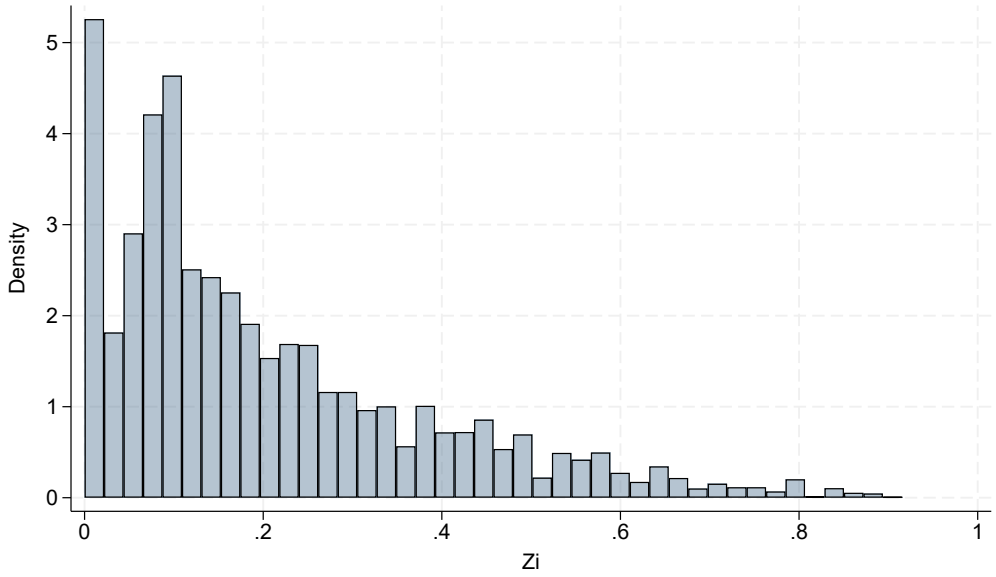
Notes: This figure shows the average cumulative number of publications published by individuals as a function of experience, separately for those who join the academic sector and those who join the private sector. The number of publications is weighted by the number of citations received up to 5 years after publication.

Figure A.7: Instrumental variable, assumptions



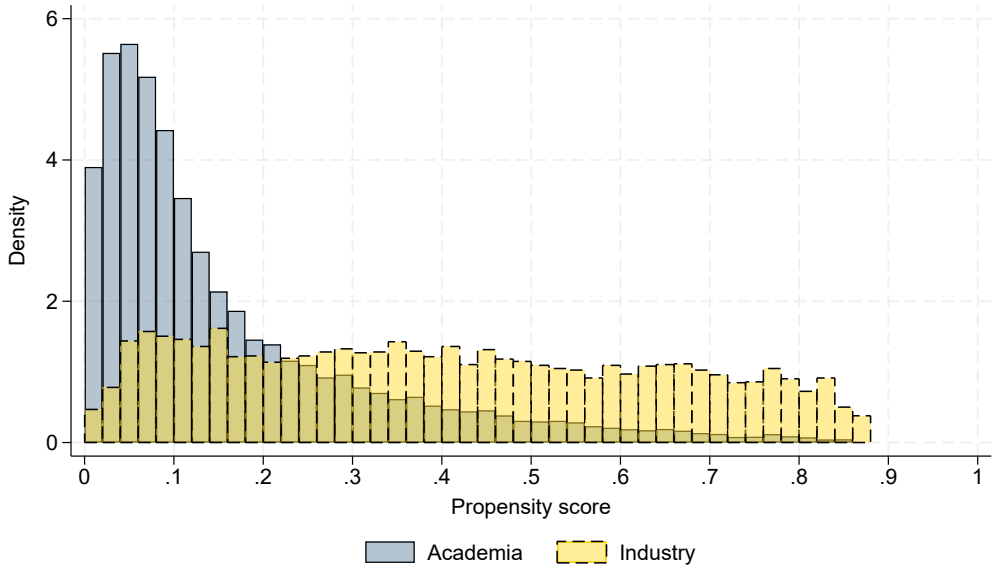
Notes: The utility function of individual i in sector j is $U_i^j = \beta_i W_i^j + \alpha_i 1\{j = \text{Industry}\}$, where α_i captures individual i's preference for Industry relative to Academia and β_i reflects sensitivity to earnings. W_i^j can be decomposed into $W_i^j = \overline{W^j} + \delta_i^j$ where $\overline{W^j}$ captures the part of earnings explained by observables and δ_i^j captures the part of earnings related to unobservable components. Similarly, the productivity of individual i in sector j can be written as $P_i^j = \overline{P^j} + v_i^j$. Individual i joins Industry iff $\beta_i(\overline{W^I} - \overline{W^A}) + \beta_i(\delta_i^I - \delta_i^A) + \alpha_i > 0$. Section F shows that $\delta_i^I - \delta_i^A$ is correlated with δ_i^I and/or δ_i^A, which creates endogeneity issues when productivity is proxied by earnings, i.e., $P_i^j = W_i^j$. As a consequence, the exogeneity assumption requires that individuals who select into Industry (higher $\delta_i^I - \delta_i^A$) do not systematically graduate in years where Industry demand is high.

Figure A.8: Distribution of the instrumental variable Z_i



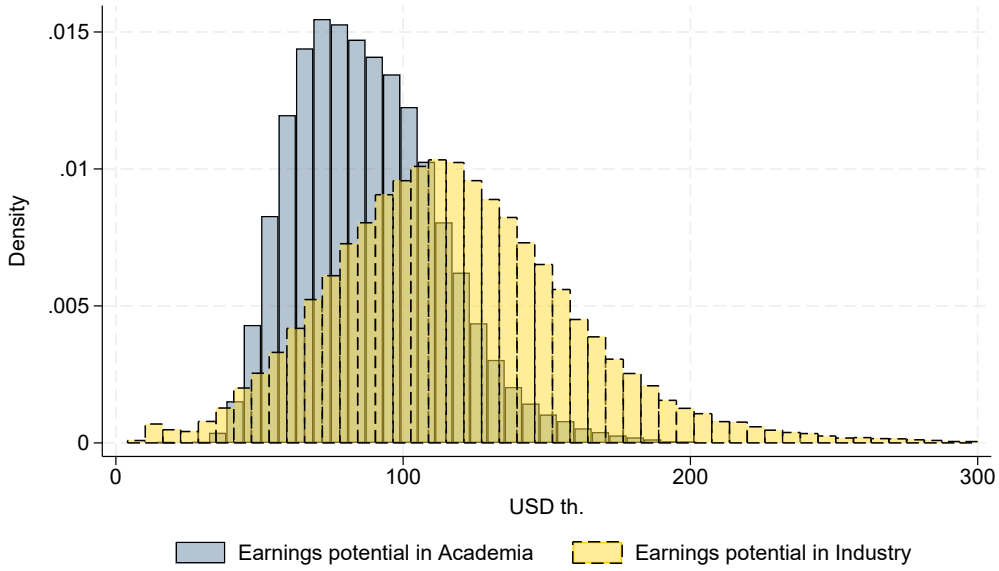
Notes: This figure shows an histogram of the instrumental variable Z_i as defined in Equation 1 for the 22,609 unique individuals in the sample. Bars with less than 5 observations are excluded for confidentiality reasons.

Figure A.9: Distribution of estimated propensity to join Industry, by sector joined



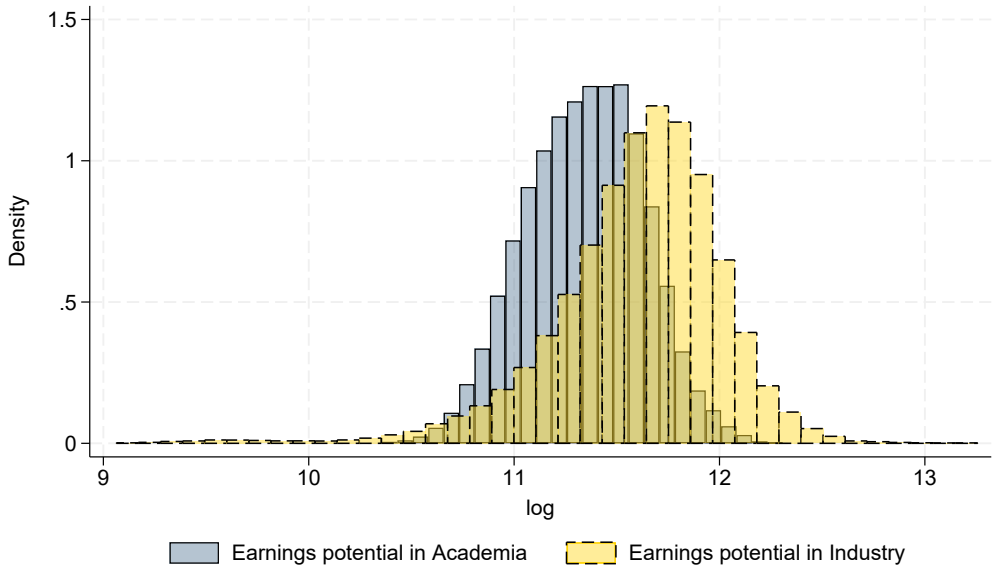
Notes: This figure displays the distribution of the estimated propensity score - the predicted probability that an individual enters Industry at graduation - based on observables (X_i) and the instrument (Z_i). The distribution is plotted separately for individuals who join Academia (blue bars with solid lines) and those who join Industry (yellow bars with dashed lines). For confidentiality, only bins with at least 10 observations are displayed.

Figure A.10: Distribution of estimated potential earnings in Academia and Industry, in level



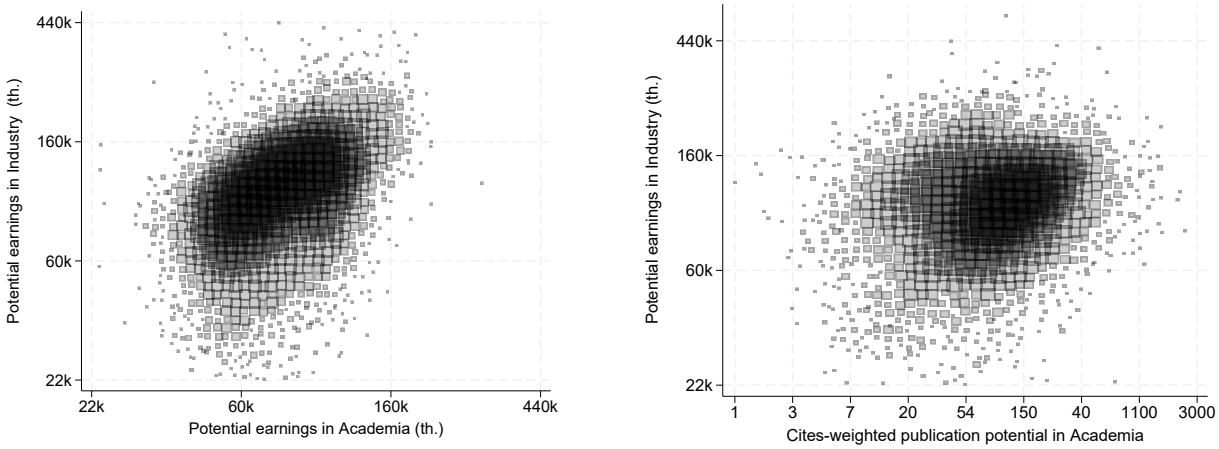
Notes: This figure shows the distribution of estimated potential earnings in Academia, transformed in 2015 USD thousands (blue bars, solid line) and estimated potential earnings in Industry, transformed in 2015 USD thousands (yellow bar, dashed line).

Figure A.11: Distribution of estimated potential earnings (log) in Academia and Industry



Notes: This figure shows the distribution of estimated potential earnings in Academia, in log (blue bars, solid line) and estimated potential earnings in Industry, in log (yellow bar, dashed line).

Figure A.12: Relationship between industrial and academic productivity

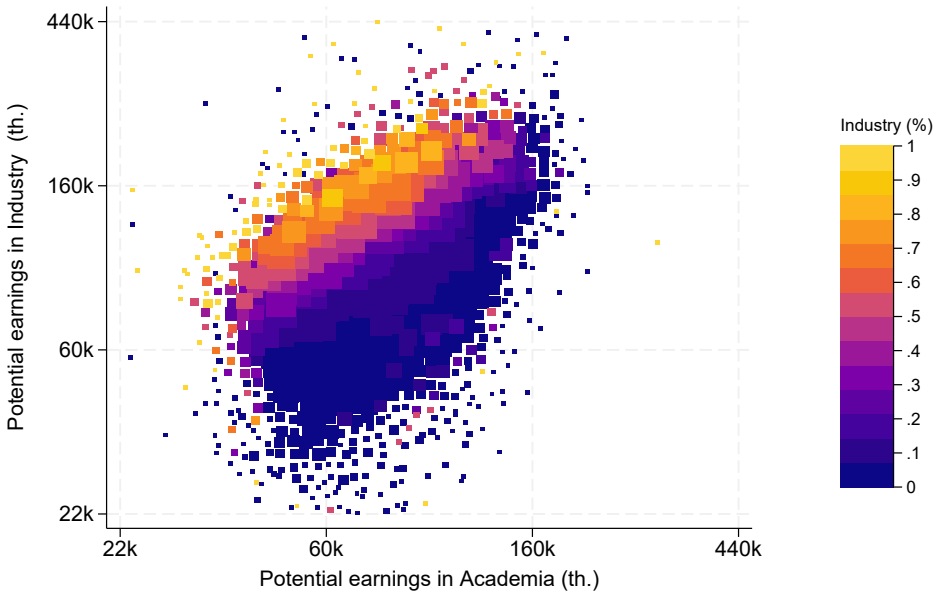


(a) Using predicted earnings

(b) Using predicted earnings and publications

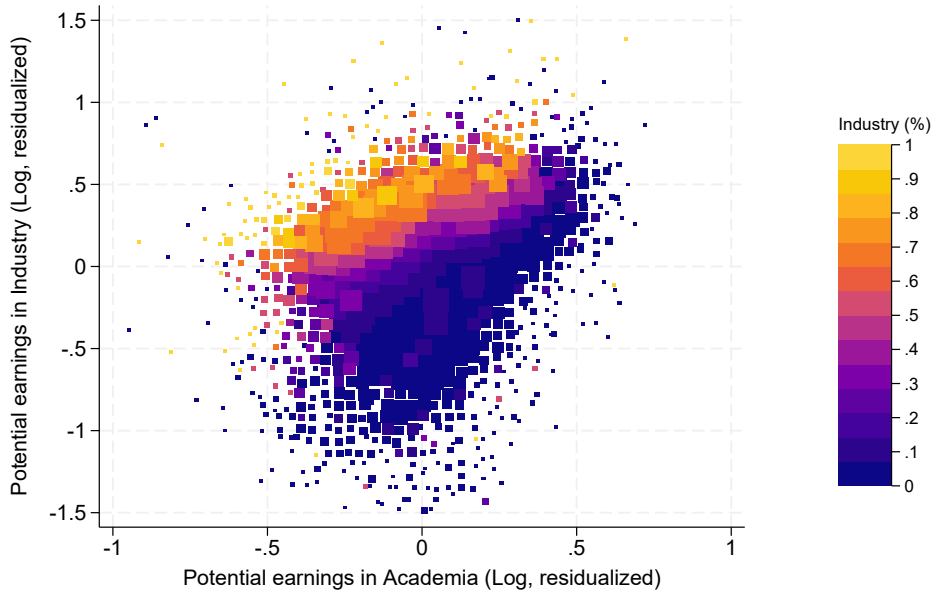
Notes: This figure shows a scatter plot of estimated productivity in Academia (x-axis) and estimated productivity in Industry (y-axis), based on Marginal Treatment Effect (MTE) estimates. Panel (a): Individuals' productivity in Academia and Industry is proxied by their predicted earnings in each sector (in thousands of 2015 USD). All earnings observations (non-residualized by experience) are kept. Panel (b): Individuals' industrial productivity is proxied by their predicted earnings in Industry (in thousands of 2015 USD), and individuals' academic productivity is proxied by their predicted publication output in Academia, weighted by citations received in the first five years.

Figure A.13: Sorting patterns by academic and industrial productivity, in level



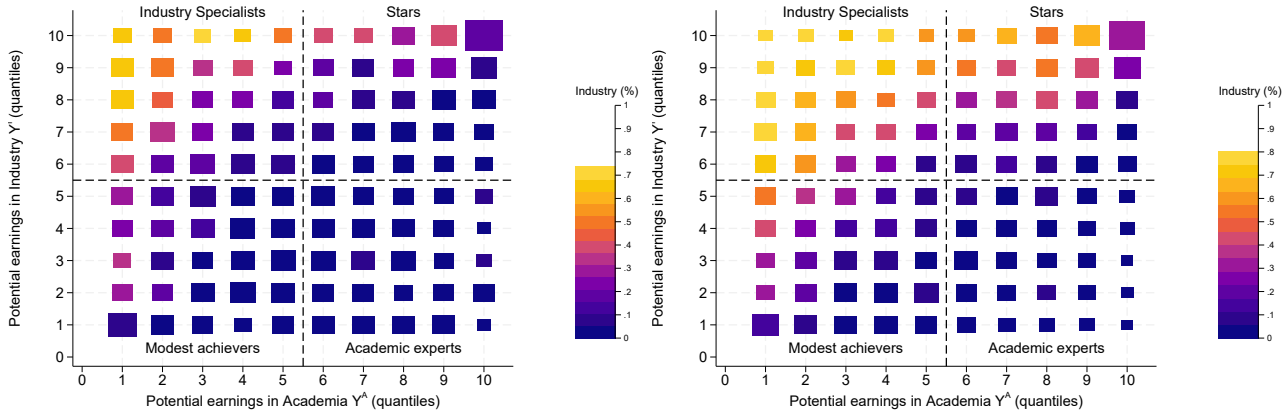
Notes: This figure shows a scatter plot of estimated potential earnings in Academia (x-axis) and estimated potential earnings in Industry (y-axis). The x-axis and y-axis are in thousands of 2015 USD. The color of the dots is proportional to the share of individuals who join the private sector (lighter=more Industry, darker=more Academia).

Figure A.14: Sorting patterns by academic and industrial productivity, (residualized)



Notes: This figure shows a scatter plot of estimated log potential earnings in Industry (y-axis) as a function of estimated log potential earnings in Academia (x-axis), residualized as in Figure 3a. The color of the dots is proportional to the share of individuals who join the private sector (lighter=more Industry, darker=more Academia).

Figure A.15: Sorting patterns by academic and industrial productivity, by gender

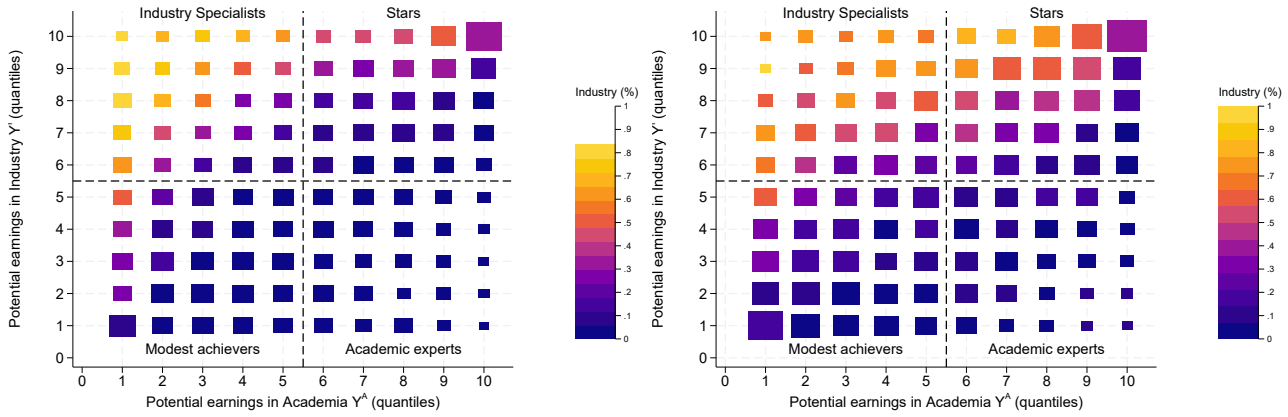


(a) Female

(b) Male

Notes: This figure shows the distribution of potential earnings and sorting outcome for females (left figure) and males (right figure). To calculate the x and y values, I first divide the sample between males and females. Within each group, I residualize potential earnings in each sector by experience and experience squared and average values at the individual level to keep one observation per individual. I then create 10 quantiles of (residualized) potential earnings in Academia (x-axis) and 10 quantiles of (residualized) potential earnings in Industry (y-axis). The size of the square is proportional to the number of individuals with the corresponding combination of x and y values. The color of the squares is proportional to the share of individuals who join the private sector.

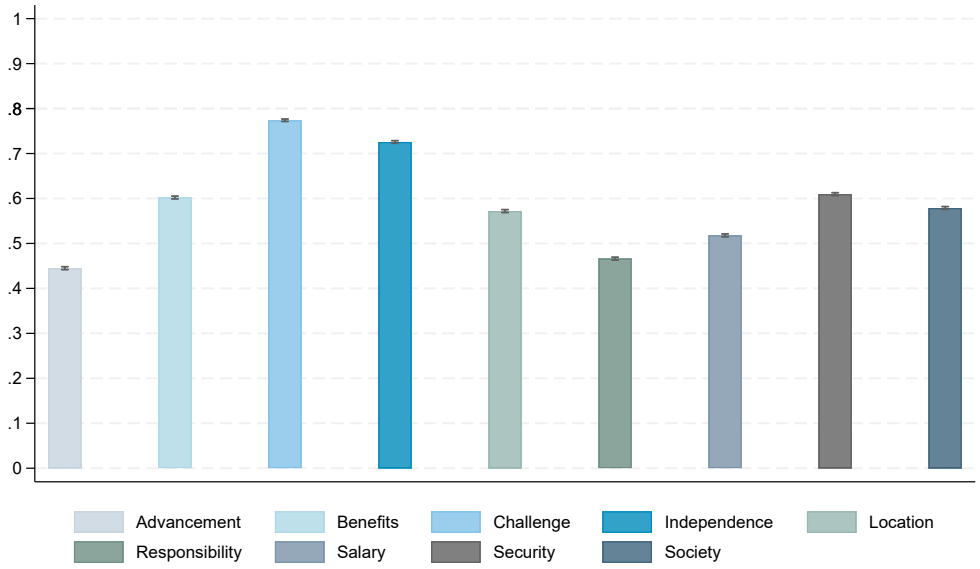
Figure A.16: Sorting patterns by academic and industrial productivity, by nationality



(a) Americans (b) Internationals

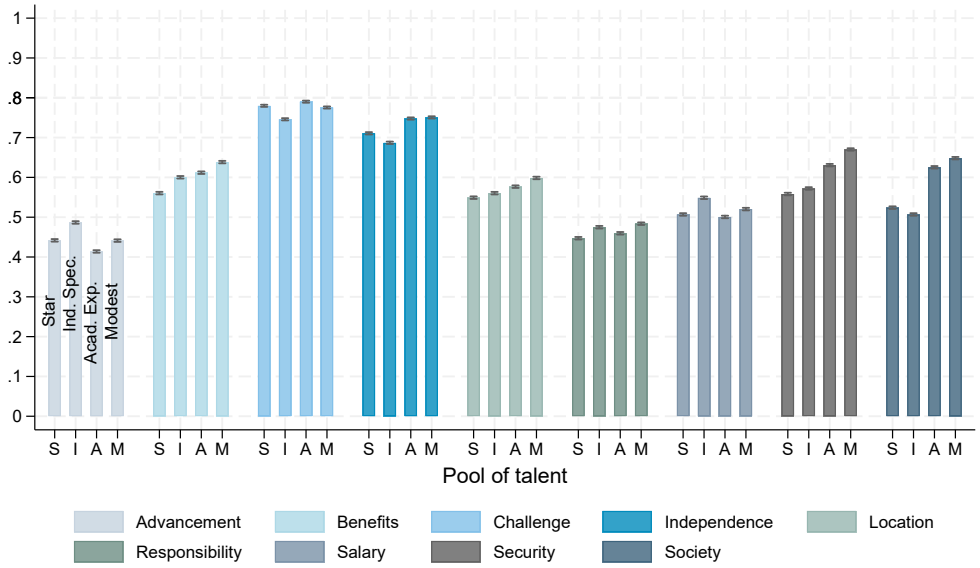
Notes: This figure shows the distribution of potential earnings and sorting outcome for Americans (left figure) and Internationals (right figure). To calculate the x and y values, I first divide the sample between Americans and Internationals. Within each group, I residualize potential earnings in each sector by experience and experience squared and average values at the individual level to keep one observation per individual. I then create 10 quantiles of (residualized) potential earnings in Academia (x-axis) and 10 quantiles of (residualized) potential earnings in Industry (y-axis). The size of the square is proportional to the number of individuals with the corresponding combination of x and y values. The color of the squares is proportional to the share of individuals who join the private sector.

Figure A.17: Preferences over job attributes



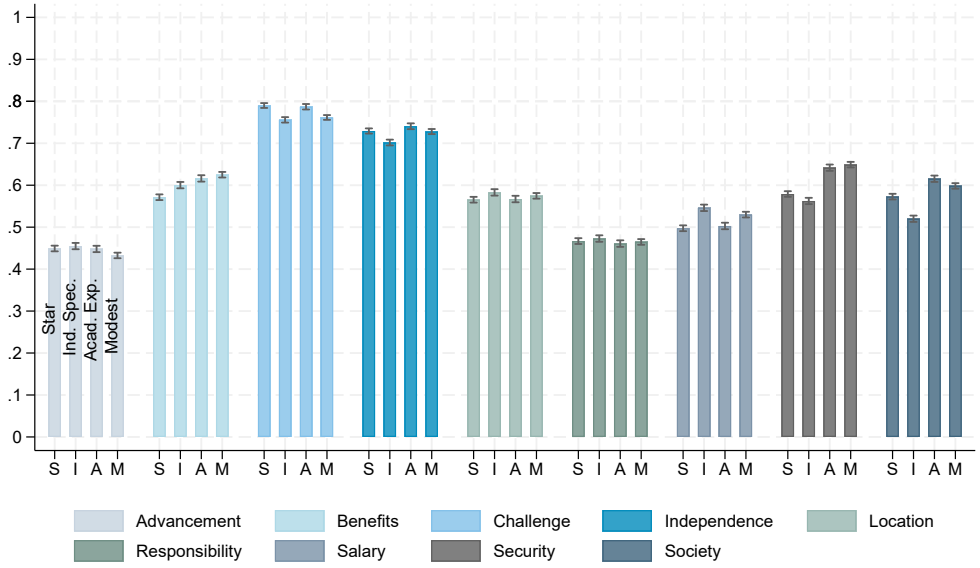
Notes: This figure shows the share of individuals who report the job attribute as 'very important'. Possible answers are 'very important', 'somewhat important', 'somewhat unimportant' or 'not important at all'.

Figure A.18: Preferences over job attributes, by profile



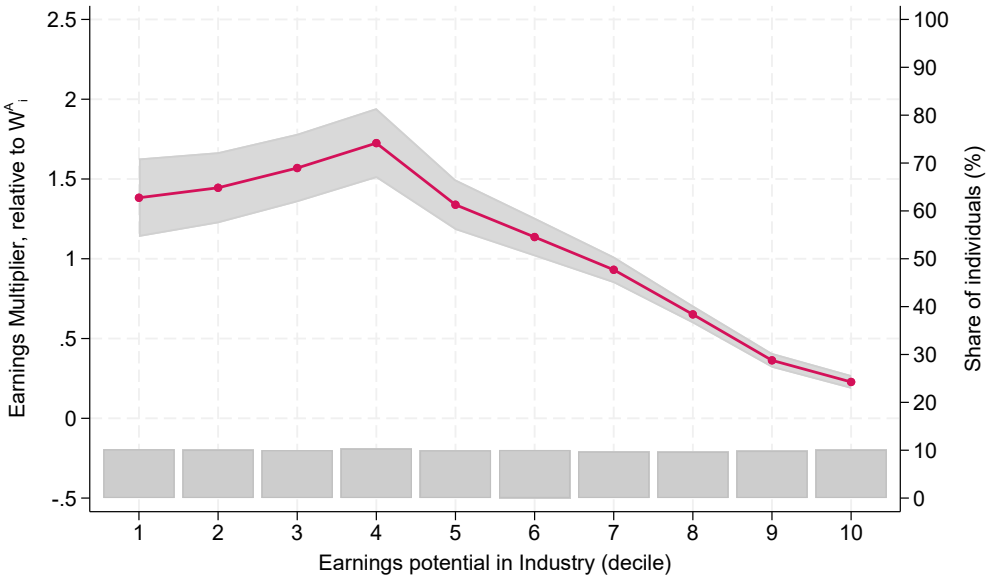
Notes: Each color represents a job attribute - from left to right: (1) job's opportunities for advancement (2) job's benefits (3) job's intellectual challenge (4) job's degree of independence (5) job's location (6) job's level of responsibility (7) job's security (8) job's salary and (9) job's contribution to society. For each attribute, the four bars represent successively: the share of 'stars' who report the job attribute as 'very important', the share of 'Industry-specialists' who report the job attribute as 'very important', the share of 'Academia-specialists' who report the job attribute as 'very important' and the share of 'modest achievers' who report the job attribute as 'very important'.

Figure A.19: Preferences over job attributes controlling for major and demographics, by profile



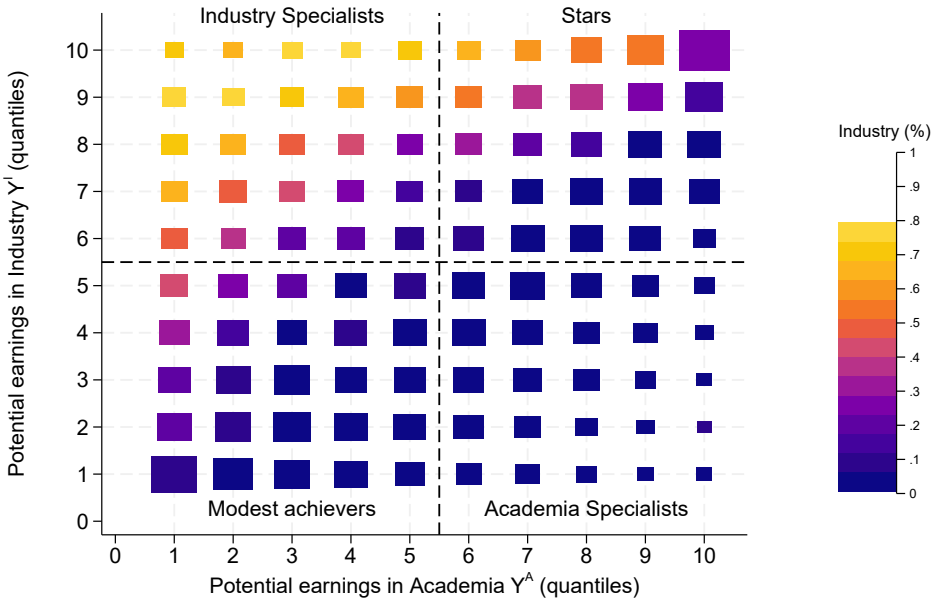
Notes: This figure shows the share of individuals who report the job attribute as 'very important', by individuals' profile, after accounting for PhD major and demographics.

Figure A.20: Estimated compensating differential, by decile of estimated industrial productivity



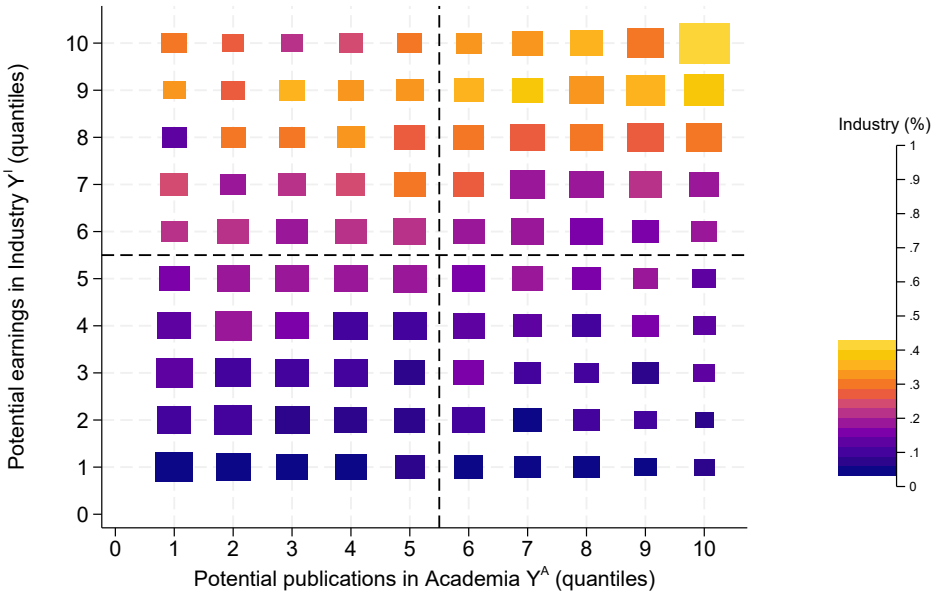
Notes: I create 10 deciles of the estimated earnings potential in Industry - which proxy for industrial productivity. For each decile, the line shows the estimated earnings premium relative to W_i^A, i.e., $\exp(\frac{-\alpha_k}{\beta}) - 1$ where α_k represents the average preference for Industry relative to Academia for individuals in decile k and β reflects the sensitivity to earnings. The shaded area indicates 95% confidence intervals, based on boostrapped standard errors clustered at the individual level. The bars represent the share of individuals in each decile.

Figure A.21: Sorting patterns by academic and industrial productivity, robustness for postdocs



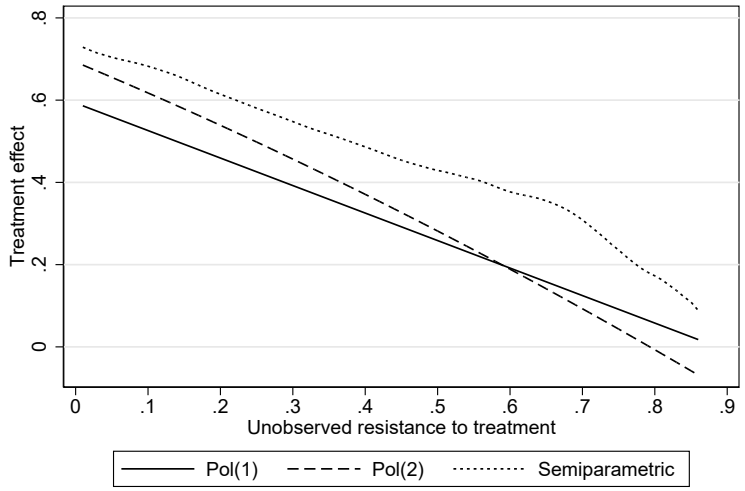
Notes: This figure estimates MTE of sector joined using as initial sector of employment for postdocs the first sector in which I observe them working 10 years post-PhD graduation (non-postdoc are still assigned as initial sector the sector in which I observe them upon PhD graduation). The x-axis shows the 10 quantiles of potential earnings in Academia and the y-axis shows the 10 quantiles of potential earnings in Industry. Potential earnings in each sector are residualized by experience and experience squared and averaged at the individual level to keep on observation per individual. The size of the square is proportional to the number of individuals with the corresponding combination of x and y values. The color of the squares is proportional to the share of individuals who join the private sector.

Figure A.22: Sorting patterns by academic and industrial productivity - excl. government



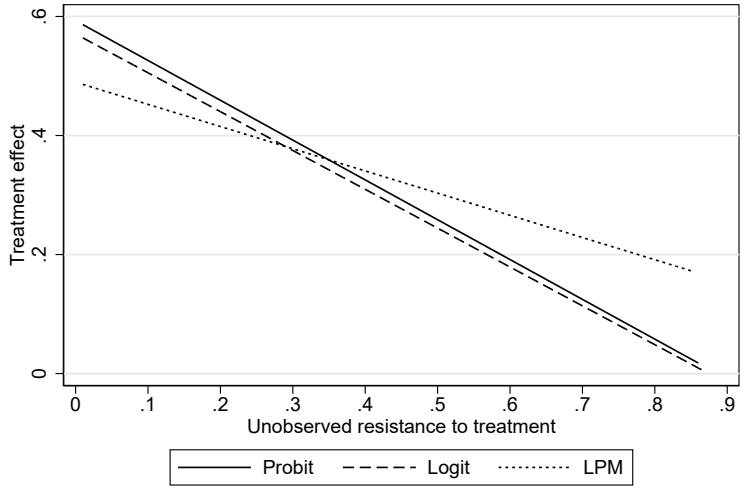
Notes: This figure excludes individuals who start in government. This figure shows the 10 quantiles of potential earnings in Academia on the x-axis and the 10 quantiles of potential earnings in Industry on the y-axis. Potential earnings in each sector are residualized by experience and experience squared and averaged at the individual level to keep on observation per individual. The size of the square is proportional to the number of individuals with the corresponding combination of x and y values. The color of the squares is proportional to the share of individuals who join the private sector.

Figure A.23: Marginal Treatment Effects curve using different functional forms



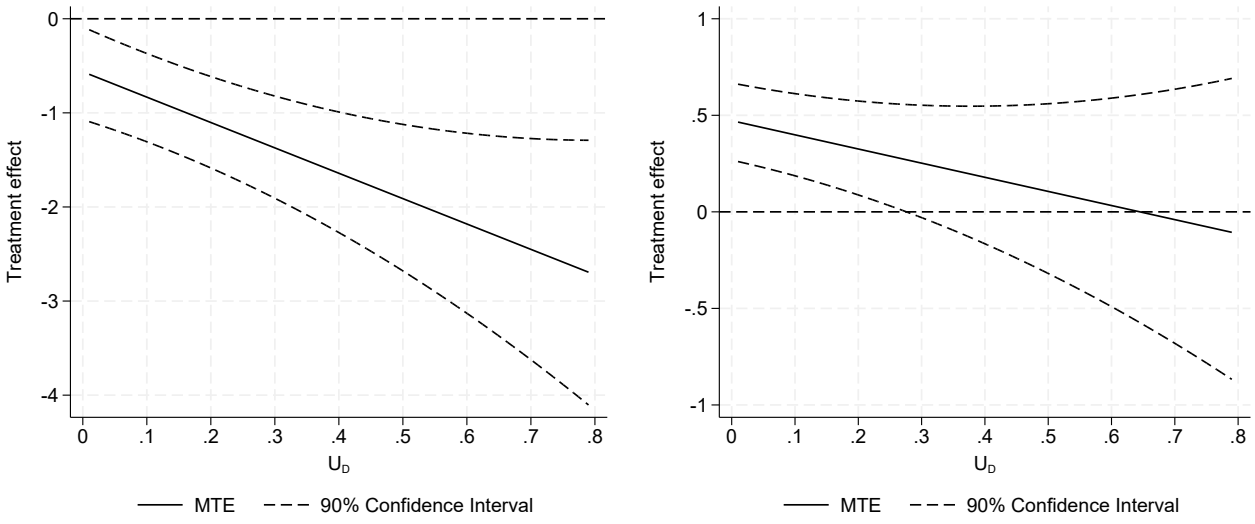
Notes: This figure shows the MTE curve using a polynomial of degree 1 (solid line), a polynomial of degree 2 (dashed line) and a semiparametric approach (dotted line).

Figure A.24: Marginal Treatment Effects curve using different first-stage specifications



Notes: This figure shows the MTE curve using a probit function (solid line), a logit function (dashed line) and a linear probability model (dotted line) for the specification of the first-stage.

Figure A.25: MTE Curves



(a) Log(Publications)

(b) Log(Earnings)

Notes: These figures show the Marginal Treatment Effects curves for the (log) publications and (log) earnings outcomes, estimated with the separate approach and a polynomial of degree 1. Standard errors are bootstrapped at the PhD cohort x PhD major level. The dotted lines represent the 90% confidence interval.

Appendix B Tables

Table B.1: OLS regression of Log(Earnings) on sector joined

	Log(Earnings)				
	(1)	(2)	(3)	(4)	(5)
Industry	0.486***	0.488***	0.413***	0.412***	0.397***
	(0.010)	(0.010)	(0.013)	(0.012)	(0.012)
Experience		Yes	Yes	Yes	Yes
PhD major FE			Yes	Yes	Yes
Doct. inst. FE				Yes	Yes
Demographics					Yes
Observations	69,917	69,917	69,917	69,917	69,917
R-sq	0.09	0.18	0.22	0.24	0.26

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The outcome is Log(Earnings). *Industry* is an indicator equal to 1 if the individual joins Industry and 0 otherwise. *Experience* includes a linear and quadratic terms for experience, calculated as the difference between survey year and PhD graduation year. *Demographics* includes an indicator equal to 1 if the individual is female, an indicator equal to 1 if the individual is white, an indicator equal to 1 if at least one of the individual's parents has a doctoral degree, an indicator for being an American citizen and birth place fixed-effects. Standard errors are clustered at the doctoral institution level.

Table B.2: First-stage

	Industry
Z_i	0.268***
	(0.043)
Demographics	Yes
Experience	Yes
PhD major FE	Yes
Doct. Inst. FE	Yes
Macro cond.	Yes
F-stat (Kleibergen-Paap)	39
F-stat (Cragg-Donald Wald)	298
Observations	69,917
R-sq	0.27

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The outcome is an indicator equal to 1 if individuals join Industry and 0 otherwise. *Demographics* includes an indicator equal to 1 if the individual is female, an indicator equal to 1 if the individual is white, an indicator equal to 1 if at least one of the individual's parents has a doctoral degree, an indicator for being an American citizen and birth place fixed-effects. *Experience* includes a linear and quadratic terms for experience, calculated as the difference between survey year and PhD graduation year. *Macro cond.* includes controls for the unemployment rate and field-specific federal funding at the time of graduation. Standard errors are clustered at the PhD major×PhD cohort level.

Table B.3: Layoff rate, by individual's profile

	1{Layoff=1}	
	(1)	(2)
Industry Specialist	0.027	
	(0.017)	
Academia Specialist	0.085**	
	(0.035)	
Modest Achiever	0.043**	
	(0.022)	
Low industrial productivity=1		0.032**
		(0.016)
PhD major FE	Yes	Yes
PhD cohort FE	Yes	Yes
Demographics	Yes	Yes
Number of survey waves FE	Yes	Yes
Observations	4,686	4,686
Mean of dep. var.	0.14	0.14

Notes: *,**,*** denote significance at 10%, 5% and 1% level respectively. The outcome is an indicator equal to 1 if the individual was ever laid off across any of the survey waves in which they appear. In column (1), the excluded category is 'star'. In column (2), 'low industrial productivity' equals 1 if the individual is either a 'modest achiever' or an 'Academia-specialist'. *Demographics* include controls for gender, race, nationality and the experience level at which I first observe each individual. Standard errors are clustered at the PhD major×PhD cohort level.

Appendix C Modeling the Correlation between Academic and Industrial Productivity

Let each individual i have three skills: k_i^G is a general skill valued by both Industry and Academia, k_i^A is a skill specific to Academia and k_i^I is a skill specific to Industry. Assume productivity is linear in skills so that:

$$P_i^A = \theta_G k_i^G + \theta_A k_i^A$$
$$P_i^B = \theta_G k_i^G + \theta_B k_i^B$$

where θ_G denotes the return to general skills, θ_A the return to Academia-specific skills and θ_B the returns to Industry-specific skills. The correlation between sector-specific productivity is given by:

$$\text{Corr}(P_i^A, P_i^B) = \frac{\text{Cov}(P_i^A, P_i^B)}{\sqrt{\text{Var}(P_i^A)} \cdot \sqrt{\text{Var}(P_i^B)}}$$

Sign. The sign of the correlation is driven by the sign of the numerator, which expands as follows:

$$\text{Cov}(P_i^A, P_i^B) = \theta_G^2 \cdot \text{Var}(k_i^G) + \theta_G\theta_B \cdot \text{Cov}(k_i^G, k_i^B) + \theta_G\theta_A \cdot \text{Cov}(k_i^G, k_i^A) + \theta_A\theta_B \cdot \text{Cov}(k_i^A, k_i^B)$$

Each term captures a distinct driver of alignment across sectors:

- $\theta_G^2 \cdot \text{Var}(k_i^G)$: Contribution of general skills to alignment (always ≥ 0)
- $\theta_G\theta_B \cdot \text{Cov}(k_i^G, k_i^B)$: Do individuals with strong general skills also tend to have strong Industry-specific skills?
- $\theta_G\theta_A \cdot \text{Cov}(k_i^G, k_i^A)$: Do individuals with strong general skills also tend to have strong Academia-specific skills?
- $\theta_A\theta_B \cdot \text{Cov}(k_i^A, k_i^B)$: Do sector-specific skills co-occur (positive) or trade off (negative)?

The first term, which reflects the contribution of general skills, is always non-negative. It pushes the correlation up when general skills are both important (high θ_G) and variable (high $\text{Var}(k_i^G)$). The second and third terms capture whether individuals with strong general skills also tend to have strong sector-specific skills. If general and sector-specific skills positively co-occur, these terms are positive and increase correlation. However, if general skills are orthogonal or inversely related to one sector's skill, these terms may reduce alignment. The final term is particularly important: if individuals tend to specialize (i.e., those who are strong in one sector's specific skills are weak in the other), then $\text{Cov}(k_i^A, k_i^B)$ is negative, and the last term pulls the correlation down. This means that even if general skills are valued, the overall correlation can be low or negative if sector-specific skills are negatively related.

Magnitude. The overall magnitude of the correlation also depends on the denominator, which captures the standard deviations of productivity in each sector. The variances are given by:

$$\text{Var}(P_i^A) = \theta_G^2 \cdot \text{Var}(k_i^G) + \theta_A^2 \cdot \text{Var}(k_i^A) + 2\theta_G\theta_A \cdot \text{Cov}(k_i^G, k_i^A)$$
$$\text{Var}(P_i^B) = \theta_G^2 \cdot \text{Var}(k_i^G) + \theta_B^2 \cdot \text{Var}(k_i^B) + 2\theta_G\theta_B \cdot \text{Cov}(k_i^G, k_i^B)$$

An increase in the denominator can dilute the correlation when the variation in each sector is driven by components that don't co-move across sectors (i.e., when the numerator does not increase at the same time).

This occurs, for instance, when productivity in each sector is heavily influenced by sector-specific skills that vary widely across individuals. In this case, the denominator increases through $\theta_A^2 \cdot \text{Var}(k_i^A)$ and $\theta_B^2 \cdot \text{Var}(k_i^B)$. If sector-specific skills are positively correlated ($\text{Cov}(k_i^A, k_i^B) > 0$), the correlation remains positive but is diluted. If they are negatively correlated ($\text{Cov}(k_i^A, k_i^B) < 0$), the correlation is pulled toward zero.

Special cases.

- If only general skills matter, then:

$$\text{Corr}(P_i^A, P_i^B) = 1$$

- If only sector-specific skills matter, then:

$$\text{Corr}(P_i^A, P_i^B) = \frac{\theta_A \theta_B \cdot \text{Cov}(k_i^A, k_i^B)}{\sqrt{\text{Var}(P_i^A)} \cdot \sqrt{\text{Var}(P_i^B)}}$$

  and the sign of the correlation depends entirely on how sector-specific skills co-vary within individuals

Appendix D Marginal Treatment Effects (MTE)

Estimating MTE allows to relate treatment effect heterogeneity to individuals' observed and unobserved propensities to join the private sector. The MTE captures the treatment effect for individuals at a given value of this propensity, enabling the estimation of a distribution of treatment effects across individuals. From the MTE, we can recover the LATE (and other key estimands such as the ATE or the ATT) by computing weighted-averages of the MTE. Importantly, I can decompose this treatment effect into potential outcomes using the separate approach (Grennan et al., 2024). Using a proxy for productivity as an outcome, this means that I can recover an empirical proxy for my main parameters of interest: P_i^A and P_i^B. The MTE can be modeled as part of the generalized Roy Model (see Andresen (2018) for a more extensive discussion).

Model – I follow the generalized Roy model and assume that each individual has two potential outcomes, one in each sector:

$$P_i^j = \overline{P^j}(X_i) + \nu_i^j, \quad j \in \{A, B\}$$

where $\overline{P^j}(X_i)$ captures the component of productivity explained by observable characteristics X_i, and ν_i^j captures unobserved determinants of productivity in sector j. The selection equation of the generalized Roy model is written as:

$$1\{j(i) = B\} = \mathbf{1}\{P(Z_i, X_i) > U_{D_i}\}$$

with $1\{j(i) = B\}$ equals 1 if individual i joins sector B (the "treatment"). $P(X_i, Z_i)$ is the propensity score, i.e., the probability of entering sector B based on observed covariates X_i and an instrument Z_i. U_{D_i} is an unobserved resistance (e.g., distaste for working in sector B) assumed to be uniformly distributed between 0 and 1. Individuals with *lower* values of U_D are *more likely* to join sector B compared to individuals with *higher* values of U_D (they have a lower distaste).

At any given value of the propensity score $p = P(X_i, Z_i)$, individuals observed in Industry must have $U_{D_i} \leq p$, while those observed in Academia must have $U_{D_i} > p$. Specifying some function for the conditional expectations of the error terms allows me to use the observed outcomes of treated and untreated individuals to estimate the following conditional expectations of P_i^A and P_i^B:

- $E[P_i^B | X_i = x, j(i) = B] = x\beta_1 + E[\nu_i^B | U_{D_i} \leq p_i] = x\beta_1 + K_1(p_i)$
- $E[P_i^A | X_i = x, j(i) = A] = x\beta_0 + E[\nu_i^A | U_{D_i} > p_i] = x\beta_0 + K_0(p_i)$

Based on the functional form assumptions chosen for $K_1(p)$ and $K_0(p)$, we can recover marginal potential outcomes, i.e., expected productivity at a specific value of unobserved resistance:

- $E[P_i^B | X_i = x, U_{D_i} = u]$
- $E[P_i^A | X_i = x, U_{D_i} = u]$

Taking the difference yields the Marginal Treatment Effect (MTE):

$$MTE(x, u) = E[P_i^B - P_i^A | X_i = x, U_{D_i} = u]$$

The MTE represents the treatment effect for individuals with observable characteristics x and latent resistance level u. It captures heterogeneity in treatment effects arising from both observed and unobserved factors.

Predictions - An advantage of the MTE estimation approach is that the estimates can be combined with the data (i.e., individuals' observables and realized sector) to derive the expected response to treatment $\mathbf{E}[P_i^B -$

$P_i^A|X_i, D_i, p_i]$ for any individual in the data with observables X, sector joined D and propensity score p. This allows me to calculate individual-level expected treatment effects and importantly potential outcomes:

$$\mathbf{E}[P_i^B|X_i, D_i, p_i] = x\beta_1 + \frac{D_i - p_i}{1 - p_i} K_1(p_i) \tag{8}$$

$$\mathbf{E}[P_i^A|X_i, D_i, p_i] = x\beta_0 + \frac{p_i - D_i}{p_i} K_0(p_i) \tag{9}$$

$$\mathbf{E}[P_i^B - P_i^A|X, D, p] = x_i(\beta_1 - \beta_0) + \frac{D_i - p_i}{p_i(1 - p_i)} K(p_i) \tag{10}$$

with $K_1(p_i), K_0(p_i), K(p_i)$ are functions of the propensity score. Equation 8 represents individual i's expected productivity in sector B and will be my empirical estimate for individual i's productivity in Industry. Equation 9 represents individual i's expected productivity in sector A and will be my empirical estimate for individual i's productivity in Academia. Equation 10 is the difference between Equation 8 and Equation 9 and represents the difference in individual i's expected productivity in sector B vs sector A.

Assumptions - The assumptions needed for the calculation of the MTE are the same as for the LATE: relevance, exclusion and monotonicity. In theory, it is possible to estimate MTE with no further assumption. However, this requires full support of the propensity score in both the treated and untreated states for all values of X which is rarely feasible (Andresen, 2018). I follow the literature and further assume additive separability between the observed and unobserved components (Carneiro et al., 2011; Cornelissen et al., 2016). While this MTE framework is more restrictive than the TSLS one, it allows me to bring more nuance into treatment heterogeneity and the pattern of selection.[38] MTE can be estimated through the local IV or the separate approach. I use the latter because it allows me to estimate potential outcomes, which are my main parameters of interest.

[38]Note that the separability assumption remains much less restrictive than a joint normal distribution of (U^A, U^B, U_D) as assumed by traditional selection models (Andresen, 2018).

Appendix E Sector Choice

I present here the full model of sector choice that allows individuals to change sector over their career. Let individuals be indexed by i and time be indexed by $t \in (0, T)$. There are 2 sectors in the economy: Industry, indexed as I and Academia, indexed as A. I assume for simplicity that I observe a unique generation of individuals so that t can also be conceptualized as years of experience. Each individual works between $t = 0$ and $t = T$ and is characterized by a vector $j(i)$ representing the sectors she works in during each time period: $j(i) = (j_0(i), j_1(i), ..., j_T(i))$. For clarity, I will use j instead of $j(i)$ in the equations that follow, but the reader should keep in mind that j is individual-specific. $j = (I, I, ..., I)$ for individuals who spend their whole career in Industry and $j = (A, A, ..., A)$ for individuals who spend their whole career in Academia, but individuals are allowed to change sector during their career. I am interested in j_0, the sector where individuals *start* their career after graduating form their PhD.

Define $W_i^{j_t}$ the earnings of individual i at time t in sector $j_t \in (I, A)$ at time t. We can decompose $W_{it}^{j_t}$ into a sector-mean and an individual-specific component:

$$W_i^{j_t} = \overline{W_t^{j_t}} + \delta_i^{j_t} \tag{11}$$

with $\overline{W_t^{j_t}}$ the average earnings in sector j_t for individuals with experience t and $\delta_i^{j_t}$ an individual-specific component. The utility $u_i^{j_t}$ of individual i at time t who works in sector j_t equals:

$$u_i^{j_t} = W_i^{j_t} + \theta_i^{j_t} + \varepsilon_i^{j_t}$$

with $\theta_i^{j_t}$ individual i's taste for sector j_t and $\varepsilon_i^{j_t}$ some noise. I model individual i's utility of *starting* her career in sector $j_0 \in (I, A)$, called $U_i^{j_0}$, as a function of the sum over her career of her expected (discounted) earnings and an error term $\theta_i^{j_0}$:

$$\begin{aligned}
U_i^{j_0} &= \sum_{t=0}^{T} \rho^t \, \mathbf{E}\left[u_i^{j_t} \mid j_0\right] \\
&= \sum_{t=0}^{T} \rho^t \, \mathbf{E}\left[W_i^{j_t} + \theta_i^{j_t} + \sigma_i^{j_t} \mid j_0\right] \\
&= \sum_{t=0}^{T} \rho^t \left(\mathbf{E}\left[W_i^{j_t} \mid j_0\right] + \mathbf{E}\left[\theta_i^{j_t} \mid j_0\right]\right) \\
&= \sum_{t=0}^{T} \rho^t \left(\mathbf{E}\left[\overline{W_t^{j_t}} + \delta_i^{j_t} \mid j_0\right] + \mathbf{E}\left[\theta_i^{j_t} \mid j_0\right]\right) \\
&= \sum_{t=0}^{T} \rho^t \Big(p(j_t = I|j_0)(\overline{W_t^I} + \delta_i^I) + p(j_t = A|j_0)(\overline{W_t^A} + \delta_i^A) + \mathbf{E}\left[\theta_i^{j_t} \mid j_0\right]\Big) \\
&= \sum_{t=0}^{T} \rho^t \Big(p(j_t = I|j_0)(\overline{W_t^I} + \delta_i^I) + (1 - p(j_t = I|j_0))(\overline{W_t^A} + \delta_i^A) + \mathbf{E}\left[\theta_i^{j_t} \mid j_0\right]\Big) \\
&= \sum_{t=0}^{T} \rho^t \Big(p(j_t = I|j_0)(\overline{W_t^I} + \delta_i^I - \overline{W_t^A} - \delta_i^A) + \overline{W_t^A} + \delta_i^A + \mathbf{E}\left[\theta_i^{j_t} \mid j_0\right]\Big)
\end{aligned} \tag{12}$$

with $p(j_t = I|j_0)$ (resp. $p(j_t = A|j_0)$) being the probability that individual i works in Industry (resp. Academia)

at time t given that i started her career in sector j_0.

Individual i *starts* her career in Industry iff :

$$\begin{aligned} j_0(i) = I \iff & U_i^I - U_i^A > 0 \\ \iff & \sum_{t=0}^{T} \rho^t \Big((p(j_t = I | j_0 = I) - p(j_t = I | j_0 = A))(\overline{W_t^I} + \delta_i^I - \overline{W_t^A} - \delta_i^A) \\ & + \mathbf{E}\left[\theta_i^{j_t} |\, j_0 = I\right] - \mathbf{E}\left[\theta_i^{j_t} |\, j_0 = A\right] \Big) > 0 \end{aligned} \tag{13}$$

Because this equation includes the difference in unobserved productivity parameters $\delta_i^I - \delta_i^A$, the sorting of individuals at graduation is not random.

Appendix F Selection on gains vs selection on level

I detail here the difference between selection on level and selection on gains and how they might be correlated with each other.

The utility function of individual i in sector j is $U_i^j = \beta_i W_i^j + \alpha_i 1\{j = \text{Industry}\}$, where α_i captures individual i's preference for Industry relative to Academia and β_i reflects sensitivity to earnings. W_i^j can be decomposed into $W_i^j = \overline{W^j} + \delta_i^j$ where $\overline{W^j}$ captures the part of earnings explained by observables and δ_i^j captures the part of earnings related to unobservable components. Individual i joins Industry iff $\beta_i(\overline{W^I} - \overline{W^A}) + \beta_i(\delta_i^I - \delta_i^A) + \alpha_i > 0$.

The first selection mechanism is selection on *gains* which refers to the value of the difference $\delta_i^I - \delta_i^A$ and influences which sector individuals choose at graduation. Individuals with higher values of $\delta_i^I - \delta_i^A$ have more to gain by going to the private sector and are thus more likely to select into that sector. The second selection mechanism refers to selection on *level* which is linked to the values of δ_i^I and δ_i^A. Consider two individuals identical on observable characteristics and with the same gains $\delta_i^I - \delta_i^A$, but assume that individual 1 has higher level values of δ_i^I and δ_i^A compared to individual 2. Because earnings are the sum of a sector-mean (common to individuals 1 and 2) and an individual-specific component (higher for individual 1), individual 1's potential earning outcomes are higher than individual 2's potential earning outcomes. While the sector selection equation is only linked to $\delta_i^I - \delta_i^A$, this difference might be correlated with the levels δ_i^I and δ_i^A except if we assume constant treatment effects. To see that, let's derive the conditions under which the difference and the levels are uncorrelated:

$$
\begin{gathered}
Cov(\delta_i^I - \delta_i^A, \delta_i^I) = 0 \\
Cov(\delta_i^I - \delta_i^A, \delta_i^A) = 0 \\
\iff \\
Var(\delta_i^I) - Cov(\delta_i^A, \delta_i^I) = 0 \\
Cov(\delta_i^I, \delta_i^A) - Var(\delta_i^A) = 0 \\
\iff \\
Cov(\delta_i^A, \delta_i^I) = Var(\delta_i^I) = Var(\delta_i^A) \\
\iff \\
\frac{Cov(\delta_i^A, \delta_i^I)}{\sqrt{Var(\delta_i^I)Var(\delta_i^A)}} = \frac{Var(\delta_i^I)}{\sqrt{Var(\delta_i^I)Var(\delta_i^A)}} = 1
\end{gathered}
$$

This implies a linear relationship between δ_i^I and δ_i^A:

$$\delta_i^I = a + b\delta_i^A$$

with a and b some constant. Plugging back into the initial conditions:

$$
\begin{gathered}
Cov(\delta_i^I - \delta_i^A, \delta_i^I) = 0 \\
Cov(\delta_i^I - \delta_i^A, \delta_i^A) = 0 \\
\Longleftrightarrow \\
Cov(a + b\delta_i^A - \delta_i^A, \delta_i^I) = 0 \\
Cov(a + b\delta_i^A - \delta_i^A, \delta_i^A) = 0 \\
\Longleftrightarrow \\
Cov(a + (b-1)\delta_i^A, \delta_i^I) = 0 \\
Cov(a + (b-1)\delta_i^A, \delta_i^A) = 0 \\
\Longleftrightarrow \\
(b-1)Cov(\delta_i^A, \delta_i^I) = 0 \\
(b-1)Var(\delta_i^A) = 0
\end{gathered}
$$

which implies that $b = 1$ or δ_i^I and δ_i^A being constant. In both cases, this implies that the difference $\delta_i^I - \delta_i^A$ is a constant.