Intellectual Mobility Frictions*

Jordan Bisset† Dennis Verhoeven‡

September 9, 2024

Abstract

In order for innovation policy to function optimally, an increase in the demand for invention must be met with a corresponding increase in supply. In this paper we document the presence of supply-side frictions which increase the cost of an inventor entering a new technological field for the first time. We build a simple model to estimate these frictions and find them to be ubiquitous across the technological space. Micro estimates show the frictions vary widely across firms and technological fields. We estimate the frictions reduces the effectiveness of innovation policy by more than 22%. At the aggregate level, our calculations suggest the frictions have increased the total cost of invention by over $330 Billion between 1990 and 2015.

Keywords: Invention, Inventor Mobility, Industrial Policy, Talent Allocation, Friction

JEL codes: O31, O38, O15

*We thank Thomas Astebro, Joep Konings, Margaret Kyle, Ralf Martin, Joel Stiebale, Otto Toivanen, Anna Valero, Reinhilde Veugelers, and Bruce Weinberg as well as participants at LSE's Prinz seminar, KU Leuven doctoral seminar, Munich Summer Institute PhD workshop, DRUID 2024, and the Competition and Innovation Summer School for their helpful comments. Verhoeven gratefully acknowledges funding from the Research Foundation Flanders (FWO grant 1241520N) and the H2020 Framework Programme (Marie Skłodowska-Curie Action DIV_INV grant 844254)

†Bocconi University, IGIER, Department of Economics| contact: jordan.bisset@unibocconi.it

‡KU Leuven, Faculty of Economics and Business; London School of Economics, Centre for Economic Performance (CEP) | contact: dennis.verhoeven@kuleuven.be.

I. Introduction

Without frictions, demand shocks – either occurring naturally through market forces or induced by policymakers – are effective in reallocating inventors between technological fields. However, factors such as retraining costs, idiosyncratic preferences, or career concerns can prevent such 'intellectual mobility'. This creates a wedge between the supply and demand of inventors at the level of technology fields. Given the recent suggestions that the pace of technological breakthrough is slowing (Gordon, 2016) and the returns to R&D are dwindling (Bloom et al., 2020), it is as important as ever to understand how inventors allocate across technological fields and the subsequent impact on the production of new inventions. This paper investigates the cost of mobility between technological fields. We argue that such cost, if it is substantial, creates a supply-side friction. This friction mediates the supply response to changes in the demand for inventors, and therefore affects the efficiency of the production of inventions.

The majority of the current empirical innovation literature seeks to inform policy through analysis of the demand side, while the role of intellectual mobility in knowledge production has received little attention. Existing research on inventor supply has focused on 'long-run supply' mechanics, such as what drives inventors into careers in science and technology (Bell et al., 2019), the importance of STEM university programs (Toivanen and Väänänen, 2016; Bianchi and Giorcelli, 2020) and encouraging the proliferation of universities that host these programs (Belenzon and Schankerman, 2013; Russell et al., 2022).[1] Labor supply to a given technological field in the short and medium term, however, depends on the willingness and ability of inventors to switch technological fields.

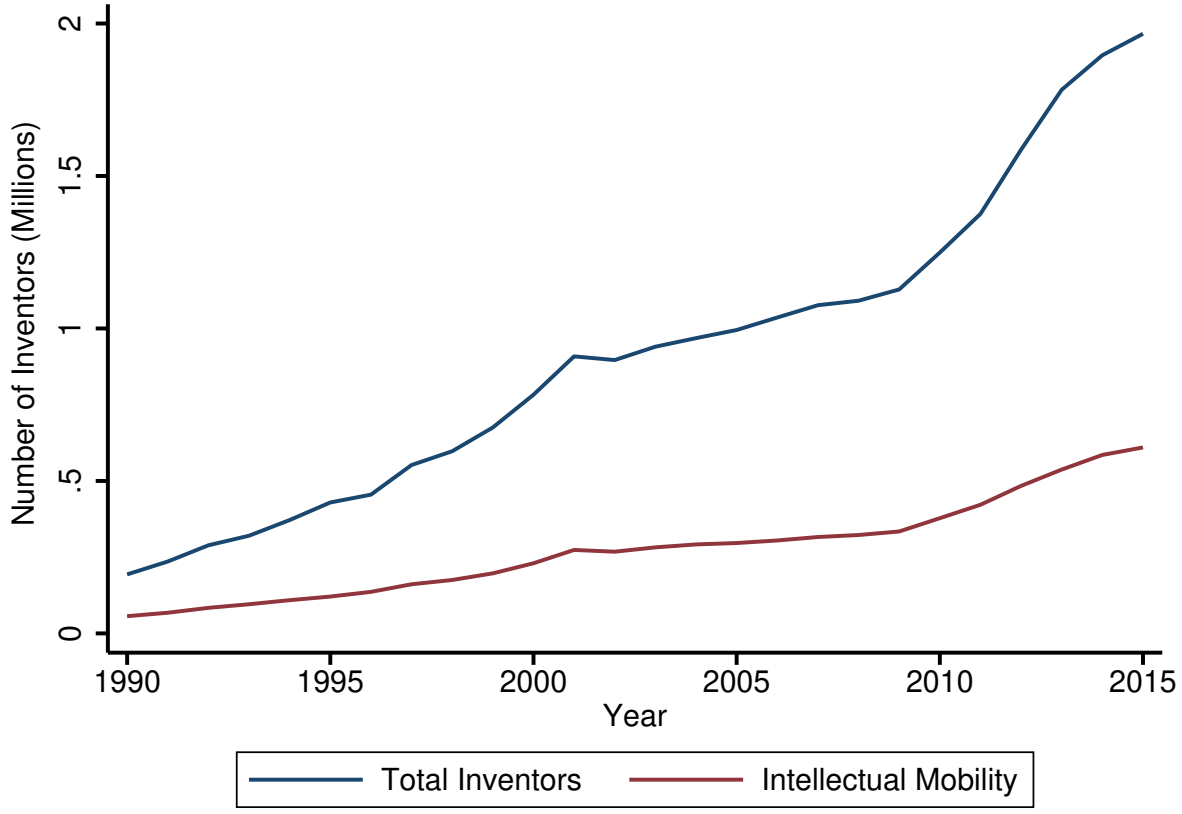
The mechanics of inventor mobility between technological fields remains a black box to the existing literature. At the aggregate level, (Jones, 2009) demonstrates that inventors have become less likely to switch between broad technological areas over time (Jones, 2009). His argument follows that in order to change field an inventor must first accrue the required knowledge to reach the technological frontier. However, reaching the ever-progressing frontier becomes more costly over time, leading to increased specialization

[1] We loosely employ the term 'long term supply' to refer to research analysing the creation of new inventors.

and reduced inventor mobility. Alternative possibilities are that inventors may be wary of entering a new technological field due to the loss of professional reputation or 'status' they hold in their existing field, or that idiosyncratic preferences have somehow evolved to be in favour of field-specific careers over time.

No matter the source, decreased mobility between aggregate fields seems to suggest the presence of some increasingly prevalent supply-side friction. The below figure 1, plots the number of active inventors, and the number of inventors switching between 5 digit CPC classes , over time. We see that even when defining technological fields at a granular level, the likelihood of intellectual mobility is decreasing over time.

Figure 1: Declining Intellectual Mobility



The figure shows trends in the number of inventors over time. The blue line shows the total number of unique inventors in our sample. The red line shows the number of inventors entering a new technological field in our sample. Technological fields are defined at the group level of the cooperative patent classification (CPC).

While such aggregate trends illustrate how supply-side frictions have evolved over time, micro-level estimates are needed to understand the implications for innovation policy. Indeed, frictions are unlikely homogeneous across the technological landscape. For instance, there is no reason to expect that switching technological fields within semiconductors entails the same cost as switching fields within biotechnology. As a result, quantifying the aggregate implications of supply-side frictions requires a framework to estimate frictions at the micro-level.

In this paper, we develop a simple framework to estimate the presence of supply-side frictions from patent data and then study the implications for innovation policy, and the aggregate cost of innovation. We develop the idea that inventors are restricted from entering new technological fields for the first time by the presence of supply-side frictions which increase the costs of innovation in the field. We posit that these frictions are widely present and highly heterogeneous across firms and technological fields.

The presence and magnitude of these frictions is of great concern to innovation policy. A friction which interferes with the supply side response can govern the efficiency of knowledge production, because the supply side response determines whether demand for invention (and hence the incentives to invest in R&D) actually translates into invention. In the absence of any supply-side frictions, inventors allocate across the knowledge space such that the production of new inventions responds to demand shocks. The higher the frictions, the larger the disconnect between innovation demand and supply. This can lead to bottlenecks such as R&D subsidies or tax breaks merely creating an additional rent for the existing inventors in the technological field, without resulting in more innovation (Goolsbee, 1998).

We begin by developing a conceptual framework where inventors are allocated across technological fields by firms who maximise their profits from innovation in these fields. In our simple setup, innovations in a given firm-field are produced by inventors from three possible sources: inventors who have previously invented within the technological field (inside inventors), inventors who have only ever invented in other technological fields (outside inventors), or inventors who have never invented before (new inventors). Our framework suggests that in a frictionless state of the world, the ratio of outside inventors to inside inventors observed in a field would be equal to the ratio of the output elasticity of outside inventors to the output elasticity of inside inventors. This is analogous for

the ratio of new inventors to inside inventors. Supply-side frictions act by lowering the concentration of outside and new inventors in the field below what would be optimal given their relative output elasticities. The more intense the friction, the larger the costs, and the greater the under-concentration of outside and new inventors in the field.

In our approach to the question, we rely on the detailed nature of the patent classification system to track inventors' activity over time and across domains of knowledge. We collect data on inventions produced by inventors, the technological field(s) an inventor is active in at each given point in time, the firm the inventor is affiliated with, and key measures of invention value, such as citations and family size. We use these measures to estimate the output elasticities of inside, outside and new inventors, and observe the relative concentration of the three inventor groups in fields over time. Spatially, the data allows us to track invention demand in technological fields at each level of the CPC hierarchy. As such, we can estimate the cost of mobility between technological fields at highly granular levels.

Based on this logic we move to our empirical setting, which at its simplest calculates the output elasticities of each inventor type based on estimating a Cobb-Douglas style innovation production function. The detailed nature of the dataset allows us to include high-dimensional fixed effects to assist with unbiased estimation. Our average friction estimates are highly significant statistically and economically. For the median firm-field-year in our data, the average frictions imply that outside inventors cost 25% more than their inside counterparts for their first project in a new field, and new inventors cost 19% more than inside inventors for their first project. The costs rise significantly when innovations are weighted based on quality.

Crucially, the data also enables us to estimate the output elasticities, and by association the supply-side frictions, at highly granular levels. We estimate heterogeneous frictions by 4 digit CPC class, and by firm by 4 digit CPC class pairs. We find that supply-side frictions are ubiquitous across the technological space and vary significantly across firms and fields. The firm-field specific estimates capture the unique production technology of individual firms and differences in innovation demand across different tech-

nological fields. Such micro-level estimates allow us to assess the impact of the estimated frictions on the efficiency of innovation policy. We conduct a simple counterfactual exercise and find that innovation policy induces 22% fewer inventions in a world with the estimated frictions than it would in a frictionless world.

The increase in the cost of aggregate innovation created by the frictions is extremely high. We calculate that the total innovation production by firms in our sample cost over $330 billion more than it would have if there were no supply-side frictions.

The remainder of the paper is structured as follows. We next place this analysis in the context of the prior art. Section II builds a simple conceptual framework which derives the supply-side frictions, and introduces our data sources. Section III details the main empirical estimations and provides the results. Section IV illustrates the implications of our results with several counterfactual exercises. Section V concludes.

Literature

The paper makes several contributions to the literature. First, we offer important evidence on the mechanics of inventor supply. As noted above, the majority of the innovation literature is focused on the demand side. Prior supply-side studies have found that increasing R&D spending can simply increase the wages to R&D workers (Goolsbee, 1998). Other work shows that increasing the total supply of inventors could be elicited through increased access to universities (Valero and Van Reenen, 2019; Belenzon and Schankerman, 2013; Russell et al., 2022) or increased access to STEM programs in particular (Toivanen and Väänänen, 2016; Bianchi and Giorcelli, 2020). Yet others illustrate that childhood conditions – e.g., whether the child is exposed to invention in particular – are vital in determining whether the child ever becomes an inventor Bell et al. (2019) (see Van Reenen et al. (2021) for a full review). We contribute to this blossoming supply-side literature by documenting a human capital friction which increases the cost of mobility between technological fields. If supply does not respond to innovation demand then innovation policy will have a sub-optimal effect on innovation. Fully understanding the presence and intensity is therefore of importance to policy-makers in optimally informing innovation policy.

We also build on the literature concerned with human capital allocation. Like the inventor supply literature, prior studies typically focus on choosing between career paths rather than directional changes within the career. Because of data limitations, many of the prior studies concerned with the allocation of talent conduct their analysis at the aggregate level (Murphy et al., 1991; Acemoglu, 1995; Philippon, 2010; Lockwood et al., 2017), not narrow fields of specialisation. In the case of inventors or scientists, recent papers have shown geographic mobility to be spurred by decisions to mitigate tax or uncertainty, or exploit R&D subsidies (Akcigit et al., 2016; Moretti and Wilson, 2017; Bisset et al., 2024). In terms of studies on mobility in the idea space, Borjas and Doran (2015) provide some evidence of mobility in the idea space of mathematics, demonstrating that American mathematicians switched fields in response to Soviet competition. Myers (2020) considers whether scientists can be incentivized to make marginal changes to their research and finds large switching costs, and Hill et al. (2021) show that researchers are less productive after changing their field of research. However, to our knowledge, we are the first to provide micro-economic evidence on the cost of inventor mobility between technological fields, and the first to connect this to implications for innovation policy.

II. Data & Model

A. Without Data

We consider a firm i that sources L inventors to technological field j in order to produce inventions $Y_{i,j}$. The inventors can have previously invented within the field ($L_{IN,i,j}$), can have previously invented but only in other fields ($L_{OUT,i,j}$), or have never previously invented ($L_{NEW,i,j}$).

The number of inventions produced in a given field j depends on the technological capability of the firm in the field ($A_{i,j}$), along with the number of inventors sourced to the field, the prior activities of the sourced inventors (denoted by subscripts $IN/OUT/NEW$), and their respective output elasticities ($\alpha_{IN}/\alpha_{OUT}/\alpha_{NEW}$).

$$Y_{i,j} = A_{i,j} L_{IN,i,j}^{\alpha_{IN}} L_{OUT,i,j}^{\alpha_{OUT}} L_{NEW,i,j}^{\alpha_{NEW}} \tag{1}$$

The firm profits from any inventions produced in the field through the revenue gained from the produced inventions net of the cost of the inventors sourced to the field. The total profits from inventions produced in the field are therefore:

$$\pi_{i,j} = \theta_{i,j} Y_{i,j} - wL_{IN,i,j} - fwL_{OUT,i,j} - swL_{NEW,i,j} \tag{2}$$

Where $\theta_{i,j}$ is a firm-field specific invention value term. It translates inventions produced by a firm in a given field into a dollar value. It captures the firm-specific gain from inventing within a particular field. w is a per invention wage that is paid to all inventors regardless of their prior activities. For inventors who have never previously invented in the field (outside or new inventors) the cost can increase due mobility costs (f) in the case of outside inventors, or start-up costs (s) in the case of new inventors. The latter can capture costs incurred by the firm in order to equip an inventor to produce inventions, such as initial training costs or on the job learning. The former, mobility costs, are dictated by intellectual mobility frictions, which can capture costs incurred through an inventor switching technological fields, such as retraining costs, a wage premium required to incentivise intellectual mobility.

Firm i selects the number of inside, outside and new inventors to be sourced to field j in order to maximise their profits from inventions in the field. This is determined by the first order conditions.

$$w = \theta_{i,j} \alpha_{IN} \frac{Y_{i,j}}{L_{IN,i,j}} \tag{3}$$

$$fw = \theta_{i,j} \alpha_{OUT} \frac{Y_{i,j}}{L_{OUT,i,j}} \tag{4}$$

$$sw = \theta_{i,j} \alpha_{NEW} \frac{Y_{i,j}}{L_{NEW,i,j}} \tag{5}$$

Substituting equation 3 into equation 5 or 4 then gives an expression for start-up frictions or intellectual mobility frictions.

$$s = \frac{\alpha_{NEW} L_{IN,i,j}}{\alpha_{IN} L_{NEW,i,j}} \tag{6}$$

$$f = \frac{\alpha_{OUT} L_{IN,i,j}}{\alpha_{IN} L_{OUT,i,j}} \tag{7}$$

The intuition of this result is simple: in the absence of any friction, the ratio of inventors of any type sourced to the field would equal the ratio of their respective output elasticities. We are particularly interested in whether any mobility costs exist due to the presence of intellectual mobility frictions. This is the case when $f > 1$. If this is the case then there are fewer outside inventors relative to inside inventors than would be expected given their relative output elasticities. The more intense the friction, the larger the disconnect between the contribution of outside inventors to innovation output and their concentration in the field.

Equation 7 suggests that the frictions can be estimated from the number of inside, outside, and new inventors the firm sourced to the field, together with estimates of their output elasticity ($\alpha_{OUT}/\alpha_{IN}/\alpha_{NEW}$). We next discuss the data before turning our attention to the empirical estimation of the frictions.

B. Data

Our analysis of the cost of intellectual mobility requires us to observe inventors and the inventions they produce over time. As our observation unit of interest is the firm-technological field pair, we must also define the firm and technological field(s) an inventor is active in throughout time. We therefore collect a large patent-level dataset from the PatentsView database of the United States Patent and Trademarks Office (USPTO) which details information on all granted utility patents at the USPTO. We link this patent-level dataset to the PATSTAT database of the European patent office. This allows us to collect information from non-USPTO patents that relate to the same invention, and more comprehensive information on the exact timing of an invention, the number of patents filed related to an invention, and the number of citations received. To identify firms in our sample we use the BVD identifier from the Orbis Global database. The Orbis Intellectual Property database links firms present in Orbis Global to the applicant field present in PATSTAT. We define the technological field (s) of each patent using the

Cooperative Patent Classification (CPC). As we are interested in understanding the cost of mobility between granularly defined technological fields, we delineate fields using the group level of the CPC hierarchy. Appendix A further details the sources consulted in constructing the sample.

The initial matched sample contains 7,094,395 unique patents produced by 2,753,597 unique inventors in 2,023,173 unique firms across 9,474 unique technological fields for a total of 98,032,751 observations between the years 1902-2021.[2]

We define a firm to be active in any technological field in which they file a patent in a given year. Each distinctive firm field combination creates a unique observation. We then isolate to fields in which a firm repeatedly patents over time by restricting the sample to firm field pairs in which at least one patent is filed each year. This gives us a sample of 4,316,698 unique firm field pairs, 4,795,297 unique patents, 2,510,172 unique inventors, 397,920 unique firms and 49,948,852 observations.

The number of inventions $Y_{i,j,t}$ that a firm i produced in a field j each year is defined by summing the number of distinct patents (y) filed by the firm in the field that year. This is illustrated by equation 8. As a patent can be assigned to multiple technological fields, a unique patent in the data can be counted as an invention in multiple firm field pairs. In addition to a simple patent count ($Y_{pat,i,j,t}$), we calculate quality-weighted measures of innovation output by summing the number of citations received by patents in each firm-field-year ($Y_{cit,i,j,t}$), and the cumulative family size ($Y_{fam,i,j,t}$) of all patents in the firm-field-year.

$$Y_{pat,i,j,t} = \sum_{y_{i,j,t}=1}^{Y_{i,j,t}} y_{i,j,t} \tag{8}$$

The distribution of unique patents across firm-field-years is shown in the below figure 2. About two thirds of firm-fields we observe have between 1 and 30 patents in a year. This varies significantly across the sample, with close to 5% of firm field observations producing more than 300 unique patents in a year.

[2]The initial sample is almost entirely confined to the 1976-2021 scope of the PatentsView database. However 0.30% of the patents are timestamped prior to 1976 as we define the patent-date at the earliest priority date at any global patent authority in order to best approximate the year of invention.

To define the number of inside, outside and new inventors sourced to each firm field year we proceed as follows: An inventor is defined as an inside inventor if they have previously been a listed inventor an a patent produced in the field in the last 10 years. An inventor is defined as an outside inventor if they have previously been listed on a patent, but not within the field in the last 10 years. An inventor who we do not observe on any prior patent is defined as a new inventor. As inventors are defined to be inside, outside or new at the level of the field, it is possible for an inventor to be considered inside in one field and outside in another field simultaneously.

For each unique patent (y) produced, an inventor (a) gives 1 unit of labor.[3]. The number of sourced inventors from inside the field is calculated by summing the number of units of labor sourced from inside inventors ($L_{IN,a}$) by the firm to the field in the year. Equation 18 illustrates the calculation for sourced inside inventors. The calculation of the labor sourced by outside and new inventors is analogous.

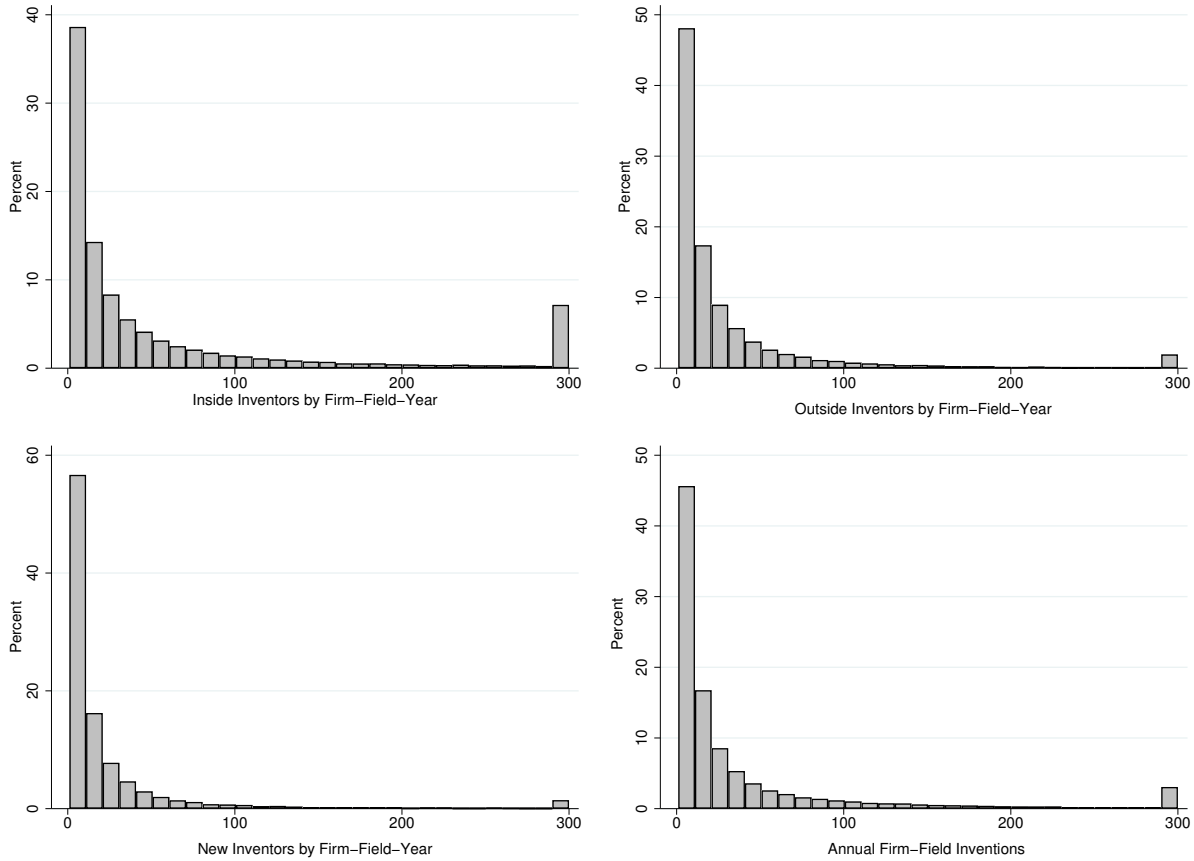
$$L_{IN,i,j,t} = \sum_{y_{i,j,t}=1}^{Y_{i,j,t}} \sum_{a_{IN,y,i,j,t}=1}^{A_{IN,y,i,j,t}} L_{IN,a,y,i,j,t} \tag{9}$$

In order to ensure that we have adequate variation in the various inventor inputs, we restrict consideration to firm field years in which at least 1 unit of inside, outside and new inventors are all utilised. This restriction reduces the sample to 2,186,265 patents, 1,036,700 inventor careers, 16,349 firms, 303,041 firm field pairs and 23,798,869 total observations between the years 1990 and 2015. The fact the latest observations in our sample (2015) occur 6 years prior to the end of the PatentsView database (2021) means we avoid the problematic patent data truncation issue (Hall et al., 2001).

The distribution of inside, outside and new inventors across firm-field-years are shown in the below figure 2. As we move from inside to outside to new inventors, the distributions become increasingly right-skewed. With 38% of firm-field-years sourcing 10 or less inside inventors, 48% sourcing 10 or less outside inventors and 57% of fields sourcing 10 or less new inventors.

[3]For example an inventor can give 2 units of labor to a firm field year if they are listed on multiple unique patents

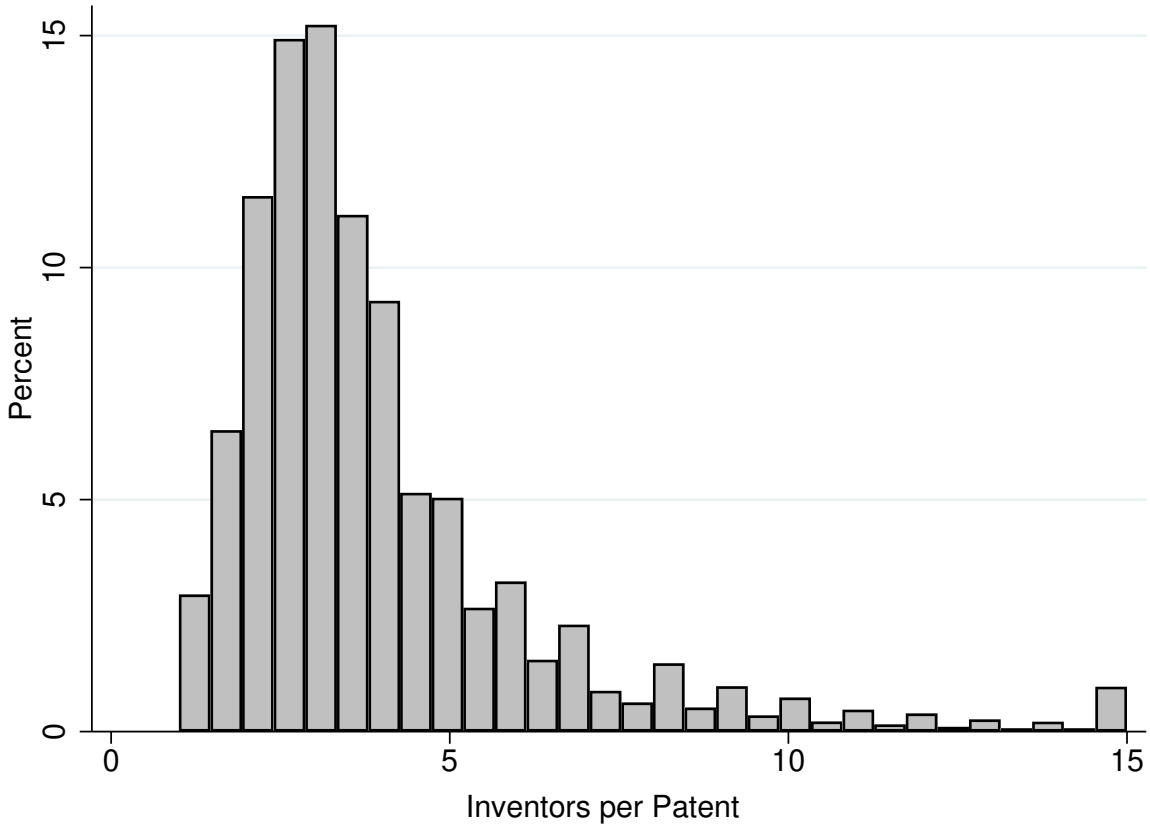
Figure 2: Firm-Field-Year distributions



The panel contains four figures which show patents, and inside, outside and new inventors are distributed across firm-field-years in the sample. In all cases observations larger than 300 are rounded down to 300 for the purpose of readability. Each figure also contains 30 bins. The upper left figure shows the distribution of the number of inside inventors sourced to a firm-field-year. The median value is 18 and the mean 92.7. The upper right figure shows the number of outside inventors sourced across firm-field-years. The median is 11 and mean 35.4. The lower left figure shows the number of new inventors active by firm-field year. The median is 8 and the mean is 27.7. The lower right figure shows the distribution of the number of unique patents per firm-field-year. The median of the distribution is 12 and the mean 46.8.

Figure 3 offers some descriptive evidence on the relationship between the inputs and output in the form of innovation production considered in equation 1. This illustrates the distribution of the total number of inventors per patent in the remaining sample. In c. 10% of inventions the number of inventors sourced is 2 or less. In more than 55% of patents, between 2 and 4 inventors were sourced. The top centile of patents using the most inventors, sourced 9 or more. Overall, in the vast majority of the sample, multiple distinct inventors are sourced in the production of a given invention, giving strong variation in the inputs to innovation production.

Figure 3: Distribution: Inventors per Invention

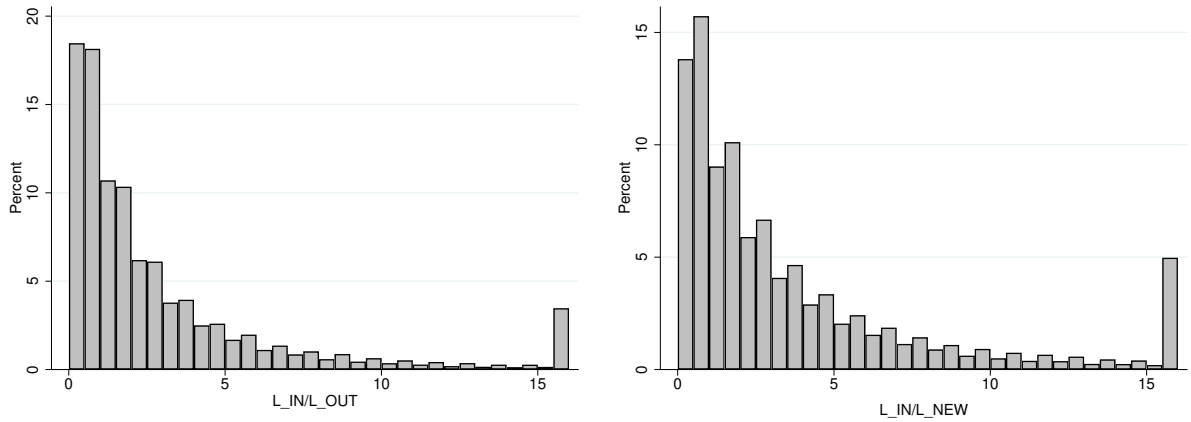


The figure shows the distribution of the total number of inventors sourced per unique patent across the sample. This is calculated at the firm-field-year level, in keeping with the unit of analysis throughout the study. The median is 4 and the mean is 5.

Key to our derivation of intellectual mobility and start-up frictions (equations 7 and 6) are the relative number of outside and new inventors sourced to the firm-field compared to the number of inside inventors. In the below panel, we provide the distributions describing the share of inside relative to new ($\frac{L_{IN,i,j,t}}{L_{NEW,i,j,t}}$) and outside inventors ($\frac{L_{IN,i,j,t}}{L_{OUT,i,j,t}}$) sourced across the firm-field-years in our sample. In both cases the majority of firm-fields source significantly more inside inventors on average. With 75% of firm-field years sourcing more inside inventors than new inventors, and roughly two thirds sourcing more inside inventors than outside inventors. This varies significantly across the sample with the top centile sourcing more than 12 times as many inside inventors as outside inventors, and nearly 16 times as many inside inventors as new inventors.

Based on equations 7 and 6, the observation that the firm-fields source more inside inventors than outside or new inventors suggests that inside inventors have a notably higher output elasticity than outside or new inventors. In order to disentangle whether the relatively high number of inside inventors is due to high mobility and start up costs, or attributable to inside inventors being particularly efficient at producing new inventions, we must empirically estimate the output elasticises of inside, outside and new inventors.

Figure 4: Participation ratios



The figures show the relative participation of inventor types across firm-field-years. The right-hand figure illustrates the ratio of inside to outside inventors. The median is 1.7 and the mean is 3.7.The left-hand figure shows the ratio of inside to new inventors. The median is 2.2 and the mean is 4.7.

III. Estimating the human capital friction

A. Output elasticity estimation

A.1 Sample Average

We now have a sample where we observe the inventors sourced and inventions produced across technological fields within firms over time. Each firm will source inventors to a technological field to produce inventions. The firm undertakes this decision with full knowledge over the output elasticity of each type of inventor, their wage, and the respective intellectual mobility and startup frictions. With equations 7 and 6 we can derive the frictions from the relative share of inventors sourced to the field-year, and from the relative output elasticities of the distinct inventor types. This requires that we can obtain unbiased estimates of the output elasticities from the observed data.

We are therefore interested in estimating the parameters of equation 10, where we allow for equation 1 to include a random disturbance to production ($\epsilon_{i,j,t}$) and a firm-field specific shock to innovation production ($\psi_{i,j,t}$).

$$Y_{i,j,t} = A_{i,j} L_{IN,i,j,t}^{\alpha_{IN}} L_{OUT,i,j,t}^{\alpha_{OUT}} L_{NEW,i,j,t}^{\alpha_{NEW}} exp(\psi_{i,j,t} + \epsilon_{i,j,t}) \tag{10}$$

We take the natural logarithm of both sides, and estimate equation 11, where $\delta_{i,j}$ are firm-field fixed effects, $\gamma_{i,t}$ are firm-year fixed effects, and $\epsilon_{i,j,t}$ is a random error term. We observe innovation output through the number of patents produced by a firm in a technological field each year ($Y_{i,j,t}$), and the number of inside, outside and new inventors ($L_{IN,i,j,t}, L_{OUT,i,j,t}, L_{NEW,i,j,t},$) sourced to each firm-field-year.[4] We do not observe the technological capability of the firm ($A_{i,j}$) or any shocks specific to the firm-field ($\psi_{i,j,t}$). The OLS estimates of the output elasticities ($\alpha_{IN}, \alpha_{OUT}, \alpha_{NEW}$) will be unbiased provided the number of inventors sourced does not vary with the error term ($\epsilon_{i,j,t}$). However, either $A_{i,j}$ or $\psi_{i,j,t}$ could influence how the firm allocates inventors across different technological fields. A firm with a strong technological capability in a given field will source more inventors to the field. Likewise, if a firm experiences a positive shock to innovation demand in one field then it may reallocate inventors from other fields to increase innovation production. Unbiased estimation of the output elasticities therefore relies on our ability to condition out these parameters.

$$log(Y_{pat,i,j,t}) = \alpha_{IN} log(L_{IN,i,j,t}) + \alpha_{OUT} log(L_{OUT,i,j,t}) + \alpha_{NEW} log(L_{NEW,i,j,t}) + \delta_{i,j} + \gamma_{i,t} + \epsilon_{i,j,t} \tag{11}$$

The inclusion of firm-field fixed effects ($\delta_{i,j}$) accounts for any firm-field specific factors which do not vary over time, such as the tacit knowledge the firm has in the field. Factors that vary over time at the firm-level but are not specific to individual technological fields are conditioned out by the firm by year fixed effects ($\gamma_{i,t}$). This accounts for firm level expansion of R&D facilities or any policy shocks, such as receipt of subsidies.[5]

[4] We also estimate versions of equation 11 using citations ($Y_{cit,i,j,t}$) and family size ($Y_{fam,i,j,t}$) to quality weight the number of patents.

[5] In theory, either capital expansion or receiving subsidies could vary over time at the firm-field level. However, given that we define the technological field at the CPC group level, it is highly unlikely that either is specific to such a granularly defined technological field.

This means we estimate our output elasticities within the firm-year. Put differently, we leverage variation in the allocation of inside, outside and new inventors, and the production of inventions between technological fields within the firm, each year.

A.2 Estimates by Technological field

In order to account for the structural differences in the production of innovations across technological fields, we relax our earlier (strong) assumptions that the output elasticities and any observed frictions are constant across the entire technological space by estimating the output elasticities individually for each technological class.

In particular, we estimate the output elasticities at the sub-class level of the CPC hierarchy. The CPC hierarchy is nested meaning that each technological field at the group level (j) nests into the broader technological field at the sub-class level. The sub-class is still a very granular definition of the bounds of a technological field. Each sub-class has on average 14 groups nested within it. [6] This gives a key advantage that the output elasticities can be heterogeneous across granularly defined technological fields while the above estimation strategy is still valid because of the variation in technological groups j within each sub-class J.

With this we estimate the output elasticities ($\alpha_{IN,J}$, $\alpha_{OUT,J}$, $\alpha_{NEW,J}$) of equation 12, where J denotes the sub-class level of the CPC hierarchy.

$$Y_{i,j,t} = A_{i,j} L_{IN,i,j,t}^{\alpha_{IN,J}} L_{OUT,i,j,t}^{\alpha_{OUT,J}} L_{NEW,i,j,t}^{\alpha_{NEW,J}} exp(\psi_{i,j,t} + \epsilon_{i,j,t}) \tag{12}$$

After taking the natural logarithm of both sides, we estimate equation 13 by OLS, and can derive field-specific frictions (f_J and s_J) through equations 14 and 15.

$$log(Y_{pat,i,j,t}) = \alpha_{IN,J} log(L_{IN,i,j,t}) + \alpha_{OUT,J} log(L_{OUT,i,j,t}) + \alpha_{NEW,J} log(L_{NEW,i,j,t}) + \delta_{i,j} + \gamma_{i,t} + \epsilon_{i,j,t} \tag{13}$$

$$s_J = \frac{\alpha_{NEW,J} L_{IN,i,J,t}}{\alpha_{IN,J} L_{NEW,i,J,t}} \tag{14}$$

[6] For example, within the technological area of chemistry, the sub-class "COMBINATORIAL CHEMISTRY; LIBRARIES, e.g. CHEMICAL LIBRARIES" nests into 9 technological groups, which include "Methods specially adapted for identifying library members" and "Methods of screening libraries"

$$f_J = \frac{\alpha_{OUT,J} L_{IN,i,J,t}}{\alpha_{IN,J} L_{OUT,i,J,t}} \tag{15}$$

While as earlier, we include firm by year fixed effects to account for any firm specific shocks which vary over time, and firm by field fixed effects, estimating the parameters at the desegregated level of the sub-class additionally benefits that any factors specific to fields at the sub-class level are additionally filtered out. Put differently the estimates condition upon factors specific to the firm-sub-class-year. This accounts for any differences in patenting practices across fields (Higham et al., 2021) and filters out any changes in technological opportunity that could influence an individual firm's innovation production decision down to the sub-class level.

A.3 Estimates by firm-field

Finally, we allow the parameters of interest to be heterogeneous by firm-sub-class, the most granular level the data allows by estimating equation 16 for each firm-sub-class pair in the sample.

$$log(Y_{pat,i,j,t}) = \alpha_{IN,i,J} log(L_{IN,i,j,t}) + \alpha_{OUT,i,J} log(L_{OUT,i,j,t}) + \alpha_{NEW,i,J} log(L_{NEW,i,j,t}) + \delta_i + \gamma_t + \epsilon_{i,j,t} \tag{16}$$

Estimating at the firm-sub-class level allows for the output elasticity, and friction, estimates to reflect the nuances in both how innovation practices vary across the technological space, and how individual firm's produce inventions. For instance the estimates can reflect different production technologies used across firms, and even different production technologies used within individual firms down to the level of 4-digit CPC classes. Aside from the capital technologies used to produce inventions, this can capture different practices of R&D managers, varying technological opportunity across the technological spaces - and how different firm's respond to such technological opportunity - along with policy shocks affecting firms, fields or the innovation activities of firm's in a given sub-class.

Such micro-level estimates also enable us to derive implications for innovation policy. For example, R&D subsidies or tax breaks are targeted towards the activities of individual firms. Therefore average estimates at the aggregate level, or at the sub-class level, would be unlikely to reflect the nuances of firm level innovation activities that is required to accurately estimate the implications that the intellectual mobility and startup frictions may have for innovation policy.

B. Average Estimates

Table 1 gives average estimates of the output elasticities and frictions across the entire sample. All columns report the estimation of equation 11. The first column with an unweighted patent count ($Y_{pat,i,j,t}$) as the dependent variable. The second column with citation weighted patents ($Y_{cit,i,j,t}$) as the dependent variable. The third column with patents weighted by family size ($Y_{fam,i,j,t}$) as the dependent variable. All columns include firm by field and firm by year fixed effects.

Across all specifications, the output elasticity of inside inventors exceeds the output elasticity of outside inventors, which exceeds the output elasticity of new inventors. Perhaps this suggests that inside inventors are able to leverage their prior experiences in the field when producing new inventions. Outside inventors likewise may benefit from their prior experience as inventors, compared to new inventors.

However, the relative output elasticity of inside inventors decreases significantly when the output of interest is quality weighted patents. As shown in columns 2 and 3 (table 1), the output elasticity ratios of both outside and new inventors ($\frac{\alpha_{OUT}}{\alpha_{IN}}$ and $\frac{\alpha_{NEW}}{\alpha_{IN}}$) are higher when either citation weighted or family size weighted patents are considered. This is particularly true in the case of outside inventors where the output elasticity increases from 0.259, when considering unweighted patents, to 0.359 when considering citation weighted patents. This reflects the findings of a significant literature on the role of inventor mobility in novel and breakthrough inventions. This suggests inventors who enter new fields are more likely to produce relatively more high value inventions, through the benefits of exposure to a new knowledge base (Arts and Fleming, 2018; Conti et al., 2014; Audia and Goncalo, 2007).

Table 1: Average Estimates

	(Patents)	(Citations)	(Family Size)
α_{OUT}	0.259***	0.359***	0.332***
	(0.001)	(0.001)	(0.001)
α_{IN}	0.347***	0.411***	0.382***
	(0.001)	(0.001)	(0.001)
α_{NEW}	0.186***	0.228***	0.240***
	(0.001)	(0.002)	(0.001)
α_{OUT}/α_{IN}	0.746	0.873	0.869
L_{IN}/L_{OUT}	1.67	1.67	1.67
f	**1.25**	**1.46**	**1.45**
α_{NEW}/α_{IN}	0.536	0.555	0.628
L_{IN}/L_{New}	2.22	2.22	2.22
s	**1.19**	**1.23**	**1.39**
Firm X Field Fixed Effects	Yes	Yes	Yes
Firm X Year FE	Yes	Yes	Yes
N	683,850	683,850	683,850

The table provides estimates of the average output elasticity of inside, outside and new inventors, and calculations of the key variables discussed in section A. All estimations take place at the firm-field level using variants of equation 11. In the first column the dependent variable is the number of patents in the firm-field-year. In the second column the dependent variable total number of citations received by patents in the firm-field-year. In the third column the dependent variable is the total family size of all patents in the firm-field-year. Standard errors are clustered at the firm-field level and reported in parenthesis. $*p < 0.1$, $**p < 0.05$, $***p < 0.01$

In accordance with equations 7 and 6, we consider the relative output elasticities in combination with the relative concentration of outside and new inventors in the firm-field. The lower panel of table 1 shows that significantly more inside inventors are sourced in the average firm-field than outside or new inventors.[7]. Based on the output elasticities it would be expected that the average field would source more inside inventors. However, the relative concentration of inside inventors far exceeds what would be expected in a state of the world with no intellectual mobility or start up frictions.

Recall, we calculate mobility frictions through the formula ($f = \frac{\alpha_{OUT}}{\alpha_{IN}} \frac{L_{IN}}{L_{OUT}}$). In a frictionless world, f would equal 1, and the observed ratio of inside to outside inventors could be exactly predicted by the ratio of their respective output elasticities. When, $f > 1$, and the human capital friction is present, there is an additional cost of mobility

[7]Firm-field data is used in calculating the ratio only if it was also used in the estimations presented in the table 1

for outside inventors entering the field for the first time. In this scenario, we would observe that, relative to inside inventors, there are fewer outside inventors in the field than would be expected based on their output elasticity. The more intense the human capital friction, the larger this discrepancy.[8]

This suggests an intellectual mobility friction ranging from 1.25 to 1.46, and a start up friction of between 1.19 and 1.39, depending on whether patents are weighted by quality in the estimation. The logic of our conceptual setting (section II) suggests that when a friction is present it shifts upwards the cost of an inventor for their first project in a new field. In the case of inventors switching fields this ranges from an additional cost of 25% to 46%. For inventors producing their first ever invention this ranges from 19% to 39%. Across all specifications, inventor mobility is afflicted with a higher friction than first time inventors. Perhaps suggesting a reluctance from inventors to engage in the retraining required to switch fields.

B.1 Robustness: Measurement Error

The average estimates in table 1 provide strong evidence of the presence of a friction which influence how inventors are allocated within firms. In a counterfactual world with no frictions, we would observe significantly more outside and new inventors in the average firm-field than we do in the data.

However, these results could be biased if there is a mismatch between the amount of inventor labor used by the firm-field and the amount we assign in the data. This is particularly problematic because we do not observe the contribution of each individual inventor to each patent. Our approach so far has been to assign 1 unit of inventor labor to the firm-field year for every inventor on the patent. However, many inventors work on multiple projects at a given time, and it seems unlikely to be true that an inventor working on ten projects at a particular point in time will offer an identical contribution to an inventor working on only 1. If this is the case, then the estimates may be afflicted with measurement error.

[8] An analogous argument holds for new inventors, based on equation 6.

To investigate, we deflate the amount of labor an inventor (a) contributes to each patent (y) by accounting for the total number of patents they are listed on in a given year. The below equation 17 illustrates this as the amount of labor contributed by inside inventor a to patent y in firm-field i, j at time t.

$$L_{IN,a,y,i,j,t} = \frac{1}{\sum_{y_{a,t}=1}^{Y_{a,t}} y_{a,t}} \tag{17}$$

We then proceed as earlier (equation 18) and sum the total amount of labor used per patent, and then sum this across all patents to calculate the amount of labor used as an input in a given firm-field-year. This is illustrated below for inside inventors.

$$L_{IN,i,j,t} = \sum_{y_{i,j,t}=1}^{Y_{i,j,t}} \sum_{a_{IN,y,i,j,t}=1}^{A_{IN,y,i,j,t}} L_{IN,a,y,i,j,t} \tag{18}$$

Under this alternative method of counting the inventor inputs, we re-estimate the average parameters using equation 11. The results are presented in table 2.

Table 2: Average Estimates - Alternative Inventor Count

	(Patents)	(Citations)	(Family Size)
α_{OUT}	0.223***	0.298***	0.282***
	(0.001)	(0.002)	(0.002)
α_{IN}	0.332***	0.380***	0.360***
	(0.001)	(0.002)	(0.002)
α_{NEW}	0.171***	0.202***	0.209***
	(0.001)	(0.003)	(0.002)
α_{OUT}/α_{IN}	0.672	0.784	0.783
L_{IN}/L_{OUT}	1.78	1.78	1.78
f	**1.20**	**1.40**	**1.39**
α_{NEW}/α_{IN}	0.515	0.532	0.581
L_{IN}/L_{New}	2.26	2.26	2.26
s	**1.16**	**1.20**	**1.31**
Firm X Field Fixed Effects	Yes	Yes	Yes
Firm X Year FE	Yes	Yes	Yes
N	193,195	193,195	193,195

The table provides a replication of the estimations in table 1 but with an alternative method of calculating the amount of labor units of inside, outside and new inventors sourced to the firm-field. Standard errors are clustered at the firm-field level and reported in parenthesis. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

The output elasticity ratios ($\frac{\alpha_{OUT}}{\alpha_{IN}}$ and $\frac{\alpha_{NEW}}{\alpha_{IN}}$) and the participation ratios ($\frac{L_{IN}}{L_{OUT}}$ and $\frac{L_{IN}}{L_{NEW}}$) remain similar to the prior estimates. The average intellectual mobility friction range falls slightly from $1.25 - 1.46$ to $1.20 - 1.40$ and the average startup friction from $1.19 - 1.39$ to $1.16 - 1.31$. However, we fail to find any major change in the statistical or economic significance of the estimates, and we still observe a significantly lower participation rate of outside and new inventors than would be expected given their relative output elasticities.

C. Results: Estimates by Technological field

We estimate equation 13 for all sub-classes for which there at least 2500 observations in the data, to ensure sufficient variation to estimate the parameters. This leaves 95 sub-classes for which we allow the output elasticities to vary. The results are visualised in the below distributions (figure 5).

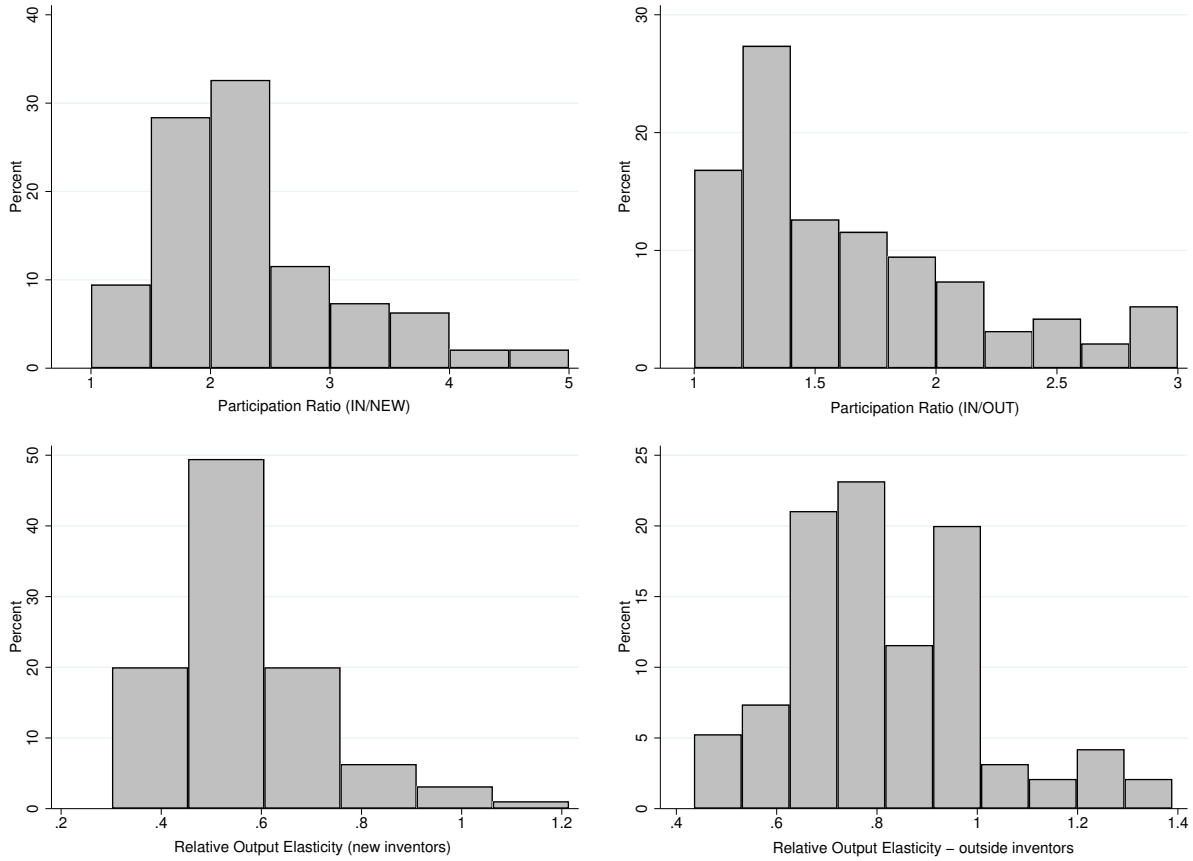
The upper half of the panel shows the participation ratios ($L_{IN,J}/L_{NEW,J}$ and $L_{IN,J}/L_{OUT,J}$) and the lower half of the panel shows the relative output elasticities ($\frac{\alpha_{OUT,J}}{\alpha_{IN,J}}$ and $\frac{\alpha_{NEW,J}}{\alpha_{IN,J}}$).

The participation value distributions are centered at 1.65 and 2.34 in the upper left and right panels, respectively, very close to the sample-wide averages provided in table 1. Even when estimating at a more granular level, we still observe very few sub-classes in which there are a higher concentration of outside or new inventors active in the firm-field-year. Only 5 of the 95 sub-classes have more outside inventors on average, and only 1 has more new inventors on average.

The relative output elasticity distributions are centered around 0.817 ($\frac{\alpha_{OUT,J}}{\alpha_{IN,J}}$) and 0.565 ($\frac{\alpha_{NEW,J}}{\alpha_{IN,J}}$). In 11 of the 95 sub-classes outside inventors have a higher output elasticity than inside inventors, 3 in the case of new inventors, suggesting that even in the absence of any intellectual mobility frictions these firm's active in these sub-classes should sourced more outside inventors.

The distributions of the intellectual mobility and startup frictions are shown in the below figure 6. The average intellectual mobility friction estimate across the 95 sub-classes is 1.27 and the average startup friction is 1.25. In 20 fields there are no start up frictions ($s < 1$) and in 8 there are no intellectual mobility frictions ($f < 1$).

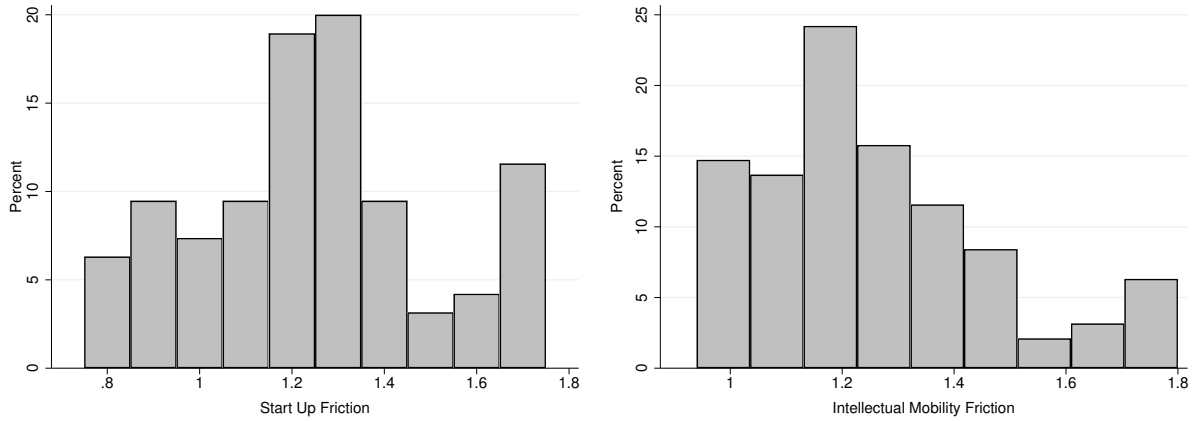
Figure 5: Results by Sub-Class: Distributions



The panel contains four figures which show the distribution of L_{IN}/L_{New}, L_{IN}/L_{OUT}, $\frac{\alpha_{OUT}}{\alpha_{IN}}$ and $\frac{\alpha_{NEW}}{\alpha_{IN}}$ based on the estimations of equation 13. The dependent variable is the natural logarithm of the number of patents in the firm-field-year (unweighted). The upper left panel shows the distribution of L_{IN}/L_{New} across sub-classes. Values less than 1 are rounded up to 1 and values more than 5 are rounded down to 5 for the purpose of readability. The upper right panel shows the distribution of L_{IN}/L_{OUT} across sub-classes. Values less than 1 are rounded up to 1 and values less than 3 are rounded down to 3 for the purpose of readability. The lower left panel shows the distribution of $\frac{\alpha_{NEW}}{\alpha_{IN}}$ across the 95 sub-class level output elasticity estimations. The lower right panel shows the distribution of $\frac{\alpha_{OUT}}{\alpha_{IN}}$ across the 95 sub-class level output elasticity estimations. In all estimations firm by year and field by firm fixed effects are included.

It would be expected that numerous fields have a friction below 1 when viewing each sub-class in isolation as the output elasticity estimates, and by association, the friction estimates reflect the specific dynamics of the fields. In some sub-classes there may be particularly high returns to inventor specialisation compared to the potentially broader knowledge base of outside inventors or new ideas of new inventors (and therefore $\alpha_{IN,J}$ relatively higher). Therefore it is striking that the intellectual mobility friction is near ubiquitous in the estimates.

Figure 6: Sub-class Friction Distributions



The left hand panel shows the distribution of the start-up friction across sub-classes. It is calculated through equation 14. Values below 0.75 are rounded up to 0.75, and values above 1.75 are rounded down to 1.75 for readability. The right hand panel shows the distribution of the intellectual mobility friction across sub-classes. It is calculated through equation 15. Values above 1.8 are rounded down to 1.8 for the purpose of readability.

Over 90% of the Sub-Classes in our sample are afflicted with a cost of mobility that leads to a lower level of inventor mobility than would be expected given their relative output elasticities. In the 20 sub-classes where the intellectual mobility friction is most intense, our estimates suggest that an outside inventor would cost 41.2-157.2% more on their first project in a field, compared to an inventor who has previously been active in the field.

D. Results: Estimates by Firm Field

Table 3 provides statistics on the distribution of the parameters obtained by estimating equation 16 for 487 firm-subclass pairs.[9]

The top panel shows the estimates when the unweighted patent count is the output of interest. The range of the estimates broadly increases compared to the estimates by sub-class indicating there is significant heterogeneity in innovation production across firms and firm-field pairs. The average friction estimates are both higher at the firm-sub-class level ($f_{i,J} = 1.34$ and $s_{i,J} = 1.32$) than at the sub-class level ($f_J = 1.27$ and $s_J = 1.25$), suggesting that the friction distributions become increasingly fat-tailed at

[9]There are more than 200,000 firm-sub-class pairs in the sample, but in order to ensure sufficient variation we only estimate equation 16 for firm-sub-classes with at least 140 unique observations. These 487 firm-subclass pairs alone account for 30% of the total innovations produced across the total sample.

more granular levels. This may be explained by firm-specific factors, such as hiring practices, or individual R&D managers being reluctant to reallocate inventors between fields in their firms. The distributions of the two frictions and the firm-sub-class level are also shown below (figure 7).

Table 3: Estimates by Firm-sub-class pairs

	Mean	SD	Median	P(10)	P(90)
$Y_{pat,i,j,t}$					
$\alpha_{OUT,i,J}/\alpha_{IN,i,J}$	0.871	0.373	0.820	0.447	1.31
$f_{i,J}$	**1.34**	**0.51**	**1.25**	**0.80**	**2.03**
$\alpha_{NEW,i,J}/\alpha_{IN,i,J}$	0.557	0.301	0.518	0.242	0.91
$s_{i,J}$	**1.32**	**0.65**	**1.22**	**0.61**	**2.19**
$Y_{cit,i,j,t}$					
$\alpha_{OUT,i,J}/\alpha_{IN,i,J}$	1.094	0.756	0.987	0.480	1.713
$f_{i,J}$	**1.69**	**1.01**	**1.50**	**0.83**	**2.67**
$\alpha_{NEW,i,J}/\alpha_{IN,i,J}$	0.616	0.627	0.502	0.112	1.181
$s_{i,J}$	**1.43**	**1.48**	**1.21**	**0.31**	**2.66**
$Y_{fam,i,j,t}$					
$\alpha_{OUT,i,J}/\alpha_{IN,i,J}$	1.080	0.653	0.959	0.471	1.721
$f_{i,J}$	**1.65**	**0.90**	**1.47**	**0.83**	**2.60**
$\alpha_{NEW,i,J}/\alpha_{IN,i,J}$	0.683	0.599	0.549	0.205	1.238
$s_{i,J}$	**1.60**	**1.44**	**1.38**	**0.49**	**2.71**
L_{IN}/L_{OUT}	1.780	0.947	1.578	0.841	2.925
L_{IN}/L_{New}	2.834	1.719	2.374	1.199	4.952

The table summarises the distribution of the key parameters obtained from estimating equation 16 across 487 firm-Sub-Class pairs. In the top panel the unweighted patent count is the dependent variable. In the middle panel citation weighted patents is the dependent variable. In the bottom panel family size weighted patents is the dependent variable. The participation ratios (L_{IN}/L_{NEW} and L_{IN}/L_{OUT}) are the same regardless of the choice of dependent variable. In all estimations field (i) and year (t) fixed effects are employed.

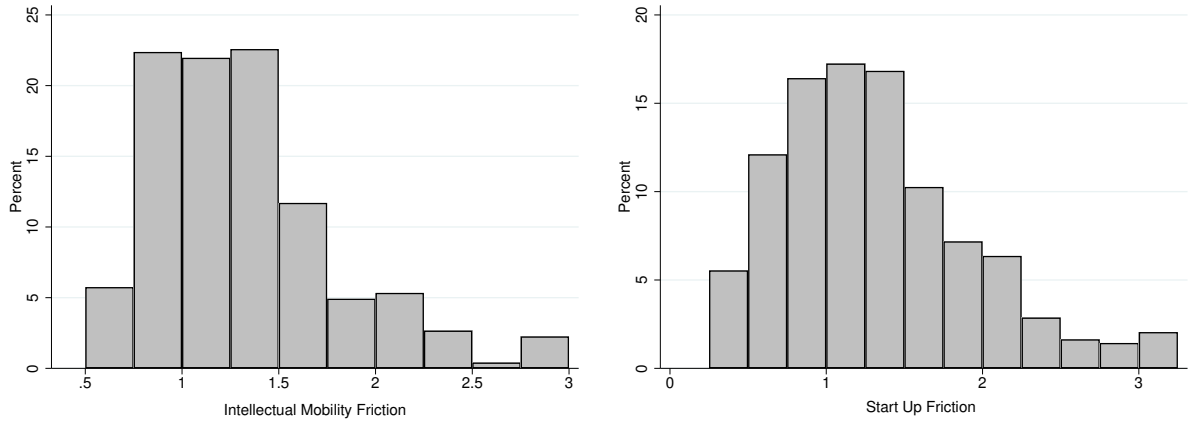
While allowing for the estimates to reflect firm and field-specific heterogeneity, the presence of frictions remains very high. 72% of firm-sub-classes are afflicted with an intellectual mobility friction, and 66% with start-up frictions. This provides strong evidence that the frictions systemically influence the sourcing of outside and new inventors across the overall technological space.

The lower two panels provide estimates when citation weighted patents ($Y_{cit,i,j,t}$) or family size weighted patents ($Y_{fam,i,j,t}$) are the output of interest. As with the average estimates (table 1), estimates at the firm-sub-class level, the relative output elasticities of outside and new inventors increase significantly when output is quality weighted. Par-

ticularly in the case of outside inventors, the relative output elasticity distribution is now centered very close to 1, suggesting that mobile inventors are particular efficient in producing high value inventions.[10] In a frictionless world, the number of outside inventors active in the average firm field would close to match the number of inside inventors, our firm-sub-class estimates suggest a significant disconnect.

We turn to exploring the implications of this disconnect for innovation policy and the aggregate cost of innovation.

Figure 7: Friction Distributions by firm-sub-class



The left hand panel shows the distribution of the intellectual mobility friction across firm-sub-classes. It is calculated as $f_{i,J} = (L_{IN}/L_{OUT}) * (\alpha_{OUT,i,J}/\alpha_{IN,i,J})$ using the estimates provided in table 3. Values below 0.5 are rounded up to 0.5, and values above 3 are rounded down to 3 for readability. The left hand panel shows the distribution of the startup friction across firm-sub-classes. It is calculated as $s_{i,J} = (L_{IN}/L_{NEW}) * (\alpha_{NEW,i,J}/\alpha_{IN,i,J})$ using the estimates provided in table 3. Values below 0.25 are rounded up to 0.25, and values above 3.25 are rounded down to 3.25 for the purpose of readability.

IV. Implications

A. Aggregate Cost

A simple and conservative way to calculate the efficiency lost due the average estimated frictions is to calculate the total cost of inventors (latter terms of equation 2) for every individual firm-field-year in a state of the world with the estimated frictions and compare to a frictionless state of the world.

With average estimated frictions ($f = 1.25$ and s=1.19):

$$TC_{friction,i,j,t} = wL_{IN,i,j,t} + 1.25wL_{OUT,i,j,t} + 1.19wL_{NEW,i,j,t} \tag{19}$$

[10] This has been suggested by a stream of papers exploring the role of inventor mobility in the production of novel and breakthrough inventions. See Arts and Fleming (2018), Conti et al. (2014) or Audia and Goncalo (2007) for reference.

With no human capital friction ($f = 1$ and $s = 1$):

$$TC_{nofriction,i,j,t} = wL_{IN,i,j,t} + wL_{OUT,i,j,t} + wL_{NEW,i,j,t} \tag{20}$$

We set the annual wage of an inventor based on the $144,096 average wage suggested for an American inventor in Jaravel et al. (2018). However, as noted earlier, inventors can simultaneously work on multiple projects in multiple different fields in a given year. To account for this we adjust the annual wage to reflect the amount of labor they contribute to each individual patent. The average inventor in our sample is involved in the production of 1.25 distinct patents each year. Therefore we calculate w ($\frac{1}{1.25} * 144,096$) as $115,924.80. This gives us the total cost of all inventors sourced each year (equation 21) and across our full quarter century sample window (equation 22).

$$\begin{cases} TC_{friction,t} = \Sigma_{i,j=1}^{303,041} TCF_{i,j,t} \\ TC_{nofriction,t} = \Sigma_{i,j=1}^{303,041} TC_{i,j,t} \end{cases} \tag{21}$$

$$\begin{cases} TC_{friction} = \Sigma_{t=1}^{26} TC_{friction,t} \\ TC_{nofriction} = \Sigma_{t=1}^{26} TC_{nofriction,t} \end{cases} \tag{22}$$

With this, we can calculate the additional costs implied by the average friction estimates for the innovation system over the last quarter century. The costs for a scenario with no frictions sum to just over $2.7 trillion, and with the estimated frictions, the total costs are just shy of $3.1 trillion. This gives a total implied cost of almost $350 billion for the average friction estimates.

B. Policy Effects

Key to our argument throughout has been that the presence of a human capital friction creates a disconnect between the supply and demand for invention. Goolsbee (1998) notes that such a disconnect can lead innovation policy to function sub-optimally. In our setting, we have identified that supply frictions are widely present across the technological space, and illustrated that this creates significant costs for the global innovation system. However, we have yet to address whether this creates to inefficiencies in innovation policy.

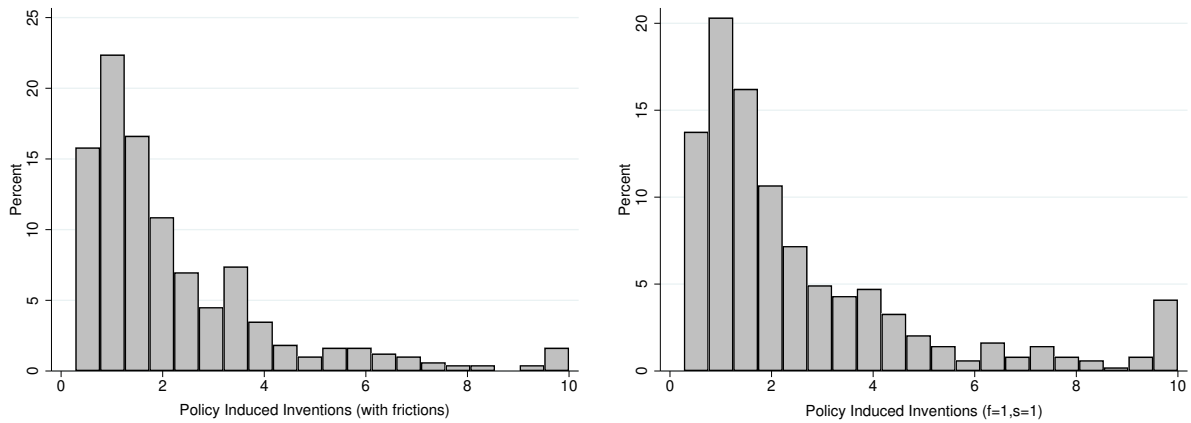
To investigate this, we use our conceptual framework (section A), in combination with the heterogeneous friction estimates (figure 7), to simulate the impact of a tax break awarded to a firm to reduce its tax burden on the revenues from invention.

We begin by estimating the average technological capability ($A_{i,j}$) and the average field specific invention value term ($\theta_{i,j}$) the firm has in the subclass.[11] This allows us to use our simple framework to calculate the number of inventors sourced, and inventions produced across different states of the world.

With this, we simulate four scenarios: (1) No policy intervention and the friction estimates shown in figure 7 (2) A 15% tax break on the activities of firm i in sub-class J, and the friction estimates shown in figure 7 (3) no policy intervention and no frictions ($f_{i,J} = 1, s_{i,J} = 1$), and (4) A 15% tax break on the activities of firm i in sub-class J, and no frictions ($f_{i,J} = 1, s_{i,J} = 1$).The tax break is simulated by increasing $\theta_{i,j}$ by 15%.

We use the four scenarios to understand how innovation policy stimulates new inventions when frictions are present, compared to the counterfactual scenario with no frictions. That is, we are interested in how many new inventions are induced by the policy when the estimated frictions are present (the difference between scenarios (1) and (2)), and how many new inventions are induced by policy when frictions are absent (the difference between scenarios (3) and (4)).

Figure 8: Policy Induced Inventions



The figures show the number of new inventions induced by a hypothetical 15% tax break on invention revenues at the firm-sub-class level. There are 479 estimates, which give the average increase in inventions for a field j year t within firm-sub-class i, J. The left figure shows the increase in inventions in a state of the world where each firm-sub-class is afflicted with their estimated frictions (figure 7). The right hand figure shows the increase in inventions in a counterfactual state of the world with no frictions. In both cases, values greater than 10 are rounded down to 10 for the purpose of readability.

[11] Equations 1 and 3 are used in conjunction with the firm-sub-class estimates, and the data.

The above figures 8 show the distribution of the average number of induced inventions per field-year i, t within the 479 firm-subclass pairs. The distributions show that the policy stimulates the production of new inventions both with and without the presence of frictions. However, the mass of the distribution shifts significantly to the right in the world without frictions. The average number of induced inventions is 2.31 in the scenario with the estimated frictions. The average number of induced inventions is 2.82 in the scenario with no frictions. This suggests that due to the presence of intellectual mobility and startup frictions, innovation policy induces 22.1% less inventions.

V. Conclusion

In order for innovation policy to function optimally, increases in demand for invention must be met by a corresponding response on the supply side. In this paper, we provide evidence of two supply-side frictions which increase the cost of sourcing new ideas to a technological field.

We illustrate the frictions through a simple a model in which firms maximize profits by selecting a number of inventors to be sourced to a technological field. The inventors can be sourced from three groups: those who have previously invented in the technological field, those who have previously only invented in other technological fields, and those who have never invented before. The frictions act by shifting the marginal cost of the latter two groups upwards relative to inside inventors.

In the absence of any friction, we would expect that the ratio of different inventors in a technological field corresponds exactly to the ratio of their output elasticities. We empirically estimate the output elasticises, and observe the number of inside, outside and new inventors in each field. With this we can estimate the fictions.

Our friction estimates suggest that there are significant costs associated with sourcing outside or new inventors to a field. Our results hold particularly large implications at the aggregate level. We calculate that the presence of our estimated friction increased the cost of the inventions observed in our sample by almost $350 Billion. A simple counterfactual exercise suggests that the estimated frictions reduces the ability of industrial policy to stimulate new inventions by more than 22%.

References

Acemoglu, D. (1995). Reward structures and the allocation of talent. *European Economic Review*, 39(1):17–33. [Cited on page 7.]

Akcigit, U., Baslandze, S., and Stantcheva, S. (2016). Taxation and the international mobility of inventors. *American Economic Review*, 106(10):2930–81. [Cited on page 7.]

Arts, S. and Fleming, L. (2018). Paradise of novelty—or loss of human capital? exploring new fields and inventive output. *Organization Science*, 29(6):1074–1092. [Cited on pages 18 and 26.]

Audia, P. G. and Goncalo, J. A. (2007). Past success and creativity over time: A study of inventors in the hard disk drive industry. *Management science*, 53(1):1–15. [Cited on pages 18 and 26.]

Belenzon, S. and Schankerman, M. (2013). Spreading the word: Geography, policy, and knowledge spillovers. *Review of Economics and Statistics*, 95(3):884–903. [Cited on pages 2 and 6.]

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713. [Cited on pages 2 and 6.]

Bianchi, N. and Giorcelli, M. (2020). Scientific education and innovation: from technical diplomas to university stem degrees. *Journal of the European Economic Association*, 18(5):2608–2646. [Cited on pages 2 and 6.]

Bisset, J., Czarnitzki, D., and Doherr, T. (2024). Inventor mobility under uncertainty. *Research Policy*, 53(1):104896. [Cited on page 7.]

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4):1104–44. [Cited on page 2.]

Borjas, G. J. and Doran, K. B. (2015). Cognitive mobility: labor market responses to supply shocks in the space of ideas. *Journal of Labor Economics*, 33(S1):S109–S145. [Cited on page 7.]

Conti, R., Gambardella, A., and Mariani, M. (2014). Learning to be edison: Inventors, organizations, and breakthrough inventions. *Organization Science*, 25(3):833–849. [Cited on pages 18 and 26.]

Goolsbee, A. (1998). Does government r&d policy mainly benefit scientists and engineers? *The American Economic Review*, 88(2):298–302. [Cited on pages 4, 6, and 27.]

Gordon, R. J. (2016). The rise and fall of american growth. In *The Rise and Fall of American Growth*. Princeton University Press. [Cited on page 2.]

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001). The nber patent citation data file: Lessons, insights and methodological tools. [Cited on pages 11 and 32.]

Higham, K., De Rassenfosse, G., and Jaffe, A. B. (2021). Patent quality: Towards a systematic framework for analysis and measurement. *Research Policy*, 50(4):104215. [Cited on page 17.]

Hill, R., Yin, Y., Stein, C., Wang, D., and Jones, B. F. (2021). Adaptability and the pivot penalty in science. *arXiv preprint arXiv:2107.06476*. [Cited on page 7.]

Jaravel, X., Petkova, N., and Bell, A. (2018). Team-specific capital and innovation. *American Economic Review*, 108(4-5):1034–73. [Cited on page 27.]

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317. [Cited on page 2.]

Lockwood, B. B., Nathanson, C. G., and Weyl, E. G. (2017). Taxation and the allocation of talent. *Journal of Political Economy*, 125(5):1635–1682. [Cited on page 7.]

Moretti, E. and Wilson, D. J. (2017). The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *American Economic Review*, 107(7):1858–1903. [Cited on page 7.]

Murphy, K. M., Shleifer, A., and Vishny, R. W. (1991). The allocation of talent: Implications for growth. *The quarterly journal of economics*, 106(2):503–530. [Cited on page 7.]

Myers, K. (2020). The elasticity of science. *American Economic Journal: Applied Economics*, 12(4):103–34. [Cited on page 7.]

Philippon, T. (2010). Financiers versus engineers: Should the financial sector be taxed or subsidized? *American Economic Journal: Macroeconomics*, 2(3):158–182. [Cited on page 7.]

Russell, L. C., Yu, L., and Andrews, M. J. (2022). Higher education and local educational attainment: Evidence from the establishment of us colleges. *The Review of Economics and Statistics*, pages 1–32. [Cited on pages 2 and 6.]

Toivanen, O. and Väänänen, L. (2016). Education and invention. *Review of Economics and Statistics*, 98(2):382–396. [Cited on pages 2 and 6.]

Trajtenberg, M., Shiff, G., and Melamed, R. (2006). The" names game": Harnessing inventors' patent data for economic research. [Cited on page 32.]

Valero, A. and Van Reenen, J. (2019). The economic impact of universities: Evidence from across the globe. *Economics of Education Review*, 68:53–67. [Cited on page 6.]

Van Reenen, J. et al. (2021). Innovation and human capital policy. Technical report, National Bureau of Economic Research. [Cited on page 6.]

Appendix

A. Data Sources - Further Details

We use the United States Patent and Trademarks Office (USPTO) PatentsView database to collect all granted utility patents filed at the USPTO between 1976 and 2020. The fact that we only observe patents that were granted prior to the end of the PatentsView database (July 7th, 2021) means there is a truncation bias towards the later years of our sample (Hall et al., 2001). To avoid any bias this may introduce, the panel only considers patents applied for prior to the end of 2015, giving all patent filings at least five years to be granted.

A key advantage of the PatentsView database is that it implements state-of-the-art name disambiguation. Name disambiguation addresses the problem that patents state inventor names rather than stable inventor identifiers. Disambiguation algorithms use the information on inventor addresses, technological classes, and co-inventors to classify different name variants as belonging to one person or, perhaps, discrete instances of the same name as being a different person. This circumvents the often discussed 'name game' problem (Trajtenberg et al., 2006) in innovation research. We use the *rawinventor* and *patent* tables from PatentsView hosted on Google BigQuery (under *patents-public-data.patentsview*) to construct a list of disambiguated inventors and their patents.

We link our set of USPTO patents to the EPO PATSTAT Global database. The key advantage of PATSTAT is that it collects information from all patent authorities globally. As such, it allows us to accurately identify the number of patents filed at global patent authorities relating to the USPTO patent detailed in the PatentsView database, and the number of citations received from patents granted at any global patent office. The PATSTAT database additionally allows us to more accurately estimate the time an invention was produced. For instance, the filing date of the USPTO patent in PatentsView may not correspond to the *first* date a patent was filed for an invention. As we are interested mainly in the timing of the *invention*, rather than *patent*, we want to timestamp an invention by its first application date.

To identify firms in our sample we use the BVD identifier from the Orbis Global database. The Orbis Intellectual Property database links firms present in Orbis Global to the applicant field present in PATSTAT. This approach circumvents the name disambiguation problem that arises from the fact that applicants in PATSTAT are not harmonized, meaning that a firm can appear in PATSTAT under different spelling variations (f.i. IBM vs. International Business Machines). We use the entity-level Orbis identifier which pertains to the legal entity registered in the various national registers that Orbis uses to compile their database. This means that a firm in our sample typically reflects an establishment in a given country.[12]

We use the Cooperative Patent Classification (CPC) to delineate technological fields. The CPC was introduced as a joint effort by the EPO and the USPTO to introduce a globally harmonized patent classification system. Patent classification is a key aspect in the legal examination process. Patents are assigned to one or more classes based on their technological content. These classes are used to search for 'prior art' – i.e. previous related patents that may diminish the inventive step claimed in the patent – when a related patent is filed. As a by-product, patent classes provide us with a taxonomy of the technological landscape. At the sub-group level – the most detailed level of the CPC – there are 252,696 classes, which nest into various hierarchical layers. We observe each level of the hierarchy for the 26 years between 1990 and 2015. There are 9 sections, 129 classes, 665 sub-classes, 9,474 groups, and 252,696 sub-groups. Any lower-level field nests into exactly one higher-level field, which nests exactly into the next higher field, and so on.

[12]We do not further consolidate entities into the larger corporations, meaning we assume that labor sourcing and R&D decisions are made by local entities rather than at the corporate level.