Searching Where Ideas are Harder to Find – The Productivity Slowdown as a Result of Firms Hindering Disruptive Innovation

Richard Bräuer*[1]

[1]Halle Institute for Economic Research

Abstract

This paper proposes to explain the productivity growth slowdown with firms consciously preventing disruptive innovation. I build an endogenous growth model with incremental and disruptive inventions and an inventor labor market where firms poach disruptive inventors to protect established technologies. I calibrate this model to the global patent landscape in 1990 and show that it predicts 52% of the decline of disruptive innovation until 2010. I confirm critical assumptions with an event study: Disruptions increase future research productivity, hurt incumbent inventors and raise the probability of future disruption. Without disruption, technology classes trend further towards incrementalism.

Keywords: general equilibrium, disruptive innovation, inventor labor market, innovation strategies, growth, general purpose technology

Classification: O12, O33, O41, J24, J42, J44

*Corresponding author: Richard Bräuer (IWH), Kleine Märkerstraße 8, 06108 Halle (Saale), Richard.braeuer@iwh-halle.de. I especially thank Eric Bartelsman and Steffen Mueller for their guidance and support. The paper has benefited greatly from the comments of Ufuk Akcigit, Jo van Biesebroeck, Sabien Dobbelaere, Michael Konig, Reinhilde Veugelers, Jonathan Deist, Matthias Mertens, Torben Klarl, Rika Stoczek and the participants of the IWH Doctoral Research Seminar, the IWH Brown Bag Seminar and the 6th AWYS Hannover. I also thank Alexander Giebler for technical assistance.

1 Introduction

The paper proposes an endogenous growth model where firms interfere with other firms' research via poaching their inventors which creates declining growth. Established firms pursue incremental innovation and are invested in specific technologies by holding a portfolio of specialized inventors. They mitigate the risk of a disruptive innovation making these technologies obsolete by hiring and benching disruptive inventors. The lack of disruptive research slows aggregate growth. This scenario is consistent with the observed data: In aggregate, firms increasingly pursue incremental research in technologies with declining returns. I perform simulation exercises to show that this effect can explain 52% of the decline in disruptive innovation observed in global patent data (PATSTAT) between 1990 and 2010.

To underpin my model, I gather stylized facts about the frequency of disruptive inventions and their repercussions. I perform a matching based event study around disruptions and find that they increase citations, patenting and the chance for a consecutive disruption, but that the effect is decaying over time. To conduct this exercise, I build upon Park, Leahey and Funk (2023) and Funk and Owen-Smith (2017) to measure patent disruptiveness via citation patterns. I construct an index of how old the patents are that citing patents reference. I.e. if a patent's citing patents do not reference older work, I deem it a disruptive innovation that spawns a new literature unconnected to the past. I apply this measure to the international patent data collected by the European Patent Office (PATSTAT) between 1980 and 2010.

I construct an endogenous growth model that reproduces my empirical findings. The actions of two types of firms drive the fate of the model economy: First, producing firms make incremental improvements of existing technologies in order to produce a product of higher quality. Second, disruptive firms do not sell any products, but try to invent a fundamentally different technology. Bill Gates and Paul Allen working in a garage to revolutionize home computing were an archetypical disruptive firm. If disruptive inventors are successful, they create a better production technology than that of any currently existing firm, but also generate large externalities: Disruptive inventions increase the productivity of future inventors and future firms, but make the current incremental inventors obsolete.

Steady technological progress requires a mixture of both types of inventions: Disruptive inventions alone never create a consumer product, only ever more advanced technologies. Incremental inventions lead to slowing technology growth over time, as incremental inventors strain against the limits of the underlying technology: Within each technology, ideas are getting harder to find. This tension between disruption and incremental growth is the central tradeoff in the model and how well the economy handles it determines economic growth.

Neither disruptive nor producing firms can conduct research on their own: Firms need inventors to make inventions for them. Firms of both types hire incremental or disruptive inventors on a search and matching labor market. Disruptive and incremental inventors enter the economy and match with firms at fixed rates. The value of each firm is determined by the stock of inventors it has hired. Incremental inventors are specialized in their current technology and cannot contribute to other technologies. Thus, whenever a firm switches the technology underlying its products, it effectively loses all incremental inventors it has hired so far. Inserting this labor market into an endogenous growth model enables the key findings with one assumption: The search and matching labor market provides both an asset intrinsically linked to technologies (incremental inventors) and a way to protect this asset from being made obsolete (by poaching disruptive inventors and stopping their work).

Successful producing firms can slow down technology disruption and overall technology growth by hiring the inventors that disruptive firms would need to innovate. This is one interpretation of the finding that being hired by large firms actually decreases inventor productivitiy (Akcigit and Goldschlag, 2023). Technological progress depends not only on investment in R&D, but also on overcoming this resistance. This sets this paper apart from the rest of the endogenous growth literature, which views innovation as the result of investment only. The longer a technology field has not been disrupted, the more of its disruptive inventors get poached, which decreases the chance for future disruption. To gauge this effect of the aging of technology fields, I estimate the key parameters of this process with the 1980-1990 portion of my data and then forecast the development of technology until 2010 using these parameters. I attribute 52% of the observed decline in disruptive innovation to the "aging" of technology fields without any model parameters changing (like the difficulty of inventions).

This paper speaks to the discussion around slowing technology growth, most notably by reconciling a set of seemingly contradictory findings: TFP growth and scientific output per researcher seem to decline, while firms hire an increasing number of researchers for non-decreasing wages (Gordon, 2016; Cowen and Southwood, 2019; Bloom et al., 2020; De Ridder, 2024). Likewise, the scientific content of patents is declining (Arora et al., 2020), despite patents with more scientific content being more valuable (Poege et al., 2019). Kalyani (2024) also argues that patents have become more derivative and have a smaller effect on firm productivity using a text-based approach. These findings are discussed as drivers of the observed slowdown in productivity growth in the literature.

Papers discussing these potential causes of the productivity slowdown in an endogenous growth setting are closest to this study. Among them, Akcigit and Ates (2023) conduct a horse race and argue that slowing technology diffusion is the most likely source of slowing technology growth, De Ridder (2024) argues for the rise of ICT technology and the resulting change in economies of scale and Olmstead-Rumsey (2019) conducts a horse race and argues that instead of ideas becoming harder to find, declining average idea quality specifically of laggards is the main driver of the decline in research productivity and the growth slowdown.

My paper explains the declining average idea quality with the decline in disruptive innovation found in publications and patents (Park, Leahey and Funk, 2023; Funk and Owen-Smith, 2017), which is again driven by the anti-competitive behavior of technology leaders. Measured idea quality decreases endogenously as firms dampen disruptive, radical innovation to reduce their risk, not because the technology process determines it. In this view, Olmstead-Rumsey (2019); Gordon (2016); Cowen and Southwood

(2019); Bloom et al. (2020); De Ridder (2024) all report troublesome trends, but these could be reversed with different policies and are not exogenously given. The proposed mechanism is self-perpetuating: The fewer disruptive inventions there are, the lower the risks for incrementally innovating firms, making them more valuable and even larger.

A fictitious social planner has to choose between incremental innovation and disruption. Which of the two he picks crucially depends on the weight that he puts on future generations: A disruptive invention will increase economic growth long-term, but the benefits will accrue to future inventors and future firms. In contrast, the current incremental inventors and producing firms unambiguously lose after a disruptive invention. If the current agents die before the growth increase from a disruptive innovation creates value, the social planner cannot compensate them and the low-growth equilibrium with incremental innovations is Pareto-optimal, even though it does not maximize GDP. If people in the model lived long enough and were patient enough, the social planner could use the additional GDP to compensate the losers from a disruptive innovation.

My model is built on the framework of Akcigit and Kerr (2018), who assume that firms are proficient in specific technology clusters. I understand technology clusters as more than just one new product, they denote distinct technologies behind multiple individual products, like "telegraphy" or "internal combustion engine". Incremental inventions within these clusters generate higher quality products. In departure from Akcigit and Kerr (2018), firms cannot invent on their own and have to hire inventors specialized in a technology cluster on a search and matching labor market. The labor market for inventors in each cluster corresponds to the results presented in the empirical chapter in Section 2.

My paper also speaks to a larger theoretical literature on market failures that misdirect innovation. Firms under-invest in research that unlocks follow-up inventions, because they cannot profit from the inventions other firms will make, as in Hopenhayn, Llobet and Mitchell (2006); Denicolò (2000); Scotchmer (1991); Acemoglu (2023). In general, firms can only appropriate a share of the overall welfare increases that result from their inventions. Since this share is not constant across inventions, firms over-invest in inventions where they can appropriate a high share of the returns (Bryan and Lemus, 2017). In the model presented here, producing firms can only fully appropriate the returns from incremental innovation, which drives their behavior and aggregate growth.

Beyond the theoretical literature, there is substantial empirical support for the monopolization of research fields, which is conceptually adjacent to the proposed model: Thompson and Kuhn (2020) use patent races between firms to compare the first and second research team and thus patent holders and followers. They find that patents preclude competitors from follow-up innovation and make the winner of patent races more dominant in the associated technology field. In the semiconductor industry, increased patent protection seems to have led to defensive patenting instead of innovation (Hall and Ziedonis, 2001). Across industries, the correlation between patent protection and innovation is negative, which Bessen and Maskin (2009) explain by the negative effect of patents on subsequent inventions. This study extends the principal insights of this literature to a context of inventor-firm labor market matching in an endogenous growth model.

This paper also links into the literature around the documented rise of firm profits and markups (Barkai, 2020; De Loecker and Eeckhout, 2017). The model predicts that firms with high market power engage in qualitatively different R&D. Only small, competitive firms invest in disruptive technology to – if successful – themselves become large firms linked to a technology. After that, their research portfolio will become much more incremental.

In a larger context, the paper relates to literature on the efficacy of the current system to reward innovative firms. The theoretical and experimental literature suggests that patents are not able to optimally steer the direction of innovation in general: If only a finite number of research direction is available, firms race each other to the most lucrative patents and incur wasteful parallel investment (Zizzo, 2002; Silipo, 2005; Breitmoser, Tan and Zizzo, 2010). Both in the US (Jaffe, 2000) and Japan (Sakakibara and Branstetter, 2001), firms do not react conclusively to substantial changes in patenting protection. Nevertheless, in my model, the market failure can be corrected by policy interventions. Since technology monopolists are misdirecting innovation, policy should break up existing monopolies and prevent mergers and buy-outs of startups. Likewise, any policy that increases the transferability of inventor skills makes technology markets larger and thus harder to monopolize.

The remainder of the paper is structured as follows: Section 2 documents stylized facts about disruptive vs. incremental innovation. Section 3 lays out the assumptions and mechanisms of the model. Section 4 simulates the model using the 1990 data to determine the quantitative importance of the model's mechanism. Section 5 discusses the response of the model to various policy interventions. Section 6 concludes the analysis.

2 Stylized Facts

2.1 Literature on Disruptive vs. Incremental Innovation

There is an active literature using firm level data to discuss the growth slowdown in developed economies. This research has generally concluded that there is a real slowdown in productivity growth, not just a measurement issue (Syverson, 2017; Reinsdorf, Byrne and Fernald, 2016; Antolin-Diaz, Drechsel and Petrella, 2017). Gordon (2016) proposed that new (impactful) ideas are getting harder and harder to find as more and more discoveries are made. He demonstrates this by estimating the worldwide researcher productivity in a series of tasks, e.g. doubling the number of transistors on a chip (Moore's Law) or crop yields per acre. Andrews, Criscuolo and Gal (2016) and Akcigit and Ates (2023) show that firm productivity dispersion has increased at the same time.

Park, Leahey and Funk (2023); Funk and Owen-Smith (2017) have documented a trend towards more incremental, less disruptive research both in publications and patents. Poege et al. (2019) show that increased incrementalism decreases the economic value of patents: Patents connected to high quality research through citations are roughly twice as valuable as other patents.

There is also substantial evidence that large firms – in contrast to small firms – lean more towards incremental improvements of existing technology (Acemoglu et al., 2016; Kerr, Nanda and Rhodes-Kropf, 2014; Kueng, Yang and Hong, 2014). The incentives that cause this behavior are also well understood theoretically (Akcigit and Kerr, 2018). When inventors get hired by these firms, their output declines (Akcigit and Goldschlag, 2023).

From these results, I draw five stylized facts that the model has to replicate:

1. Aggregate productivity growth is slowing down.

2. Researcher productivity measured for specific targets is declining.

3. Firms' research is becoming more incremental.

4. Large firms' research is more incremental than small firms' research.

5. Inventors' productivity declines when moving to large firms.

2.2 The PATSTAT data

Patent data from across the world gathered in the PATSTAT database forms the basis of my empirical strategy. This data contains the filing date of any patent application, a description of the technology and the names of firms and inventors involved. The EPO mostly relies on partner patent offices for digitization, so both coverage and the available variables vary by country. For some participating countries, the data starts in 1850, however, coverage pre-WW2 is generally low. Patents from some countries are only available from a later date onwards: E.g., Japan enters the database in the mid-seventies. Around the same time, coverage rates improve in general and the data can give a reliable picture of worldwide patent activity. I start my analysis in 1980, when the data from the major patent offices contains citations and coverage is satisfactory. Figure 1 shows the number of patents over time for selected countries. Note that the stable or shrinking number of national patents for EU countries is offset by a large increase in EU-wide EPO applications.

Figure 1: *Overview over PATSTAT*

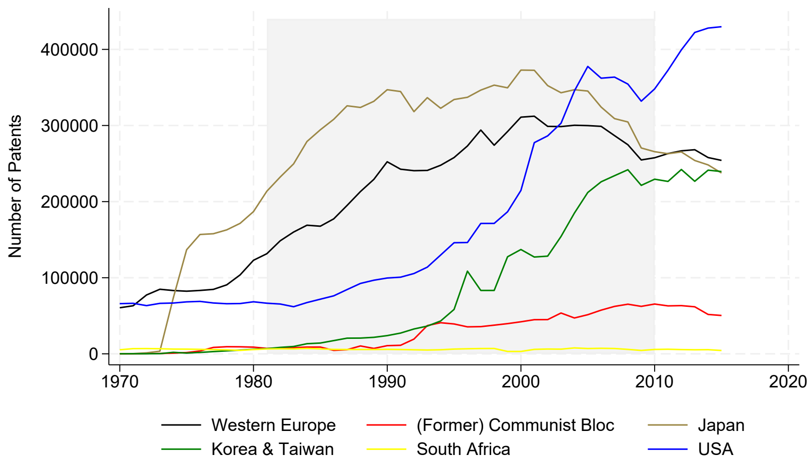

Notes: Number of patents in PATSTAT per region. The gray region marks the time period of data used in the event study in Section 2.5.
Sources: PATSTAT (European Patent Office).

Importantly, PATSTAT does not contain unique firm or person identifiers. Instead, it contains a character string written into the fields "inventor" and "applicant" on the patent. I create inventor IDs building on a large literature on name spelling unification and name disambiguation (see Appendix A and Magerman, Van Looy and Song (2006); Toole, Jones and Madhavan (2021); Li et al. (2019) for a discussion of this issue).

PATSTAT contains detailed descriptions of patents' content: Besides titles, abstracts and patent texts, the EPO assigns one or more harmonized 8 digit IPC classes to every patent. I use these IPC classes as analogues to technology fields in the theoretical section of the paper. The EPO also groups patents for the same inventions together as families and provides citations between patent families. I use these patent family citations to determine how incremental or disruptive any given patent is and aggregate these measures to the IPC-class level.

2.3 Measuring patent "disruptiveness"

To determine how disruptive specific technologies are, I follow the general strategy of Park, Leahey and Funk (2023) and Funk and Owen-Smith (2017) and look at the citation patterns around patent p. Both papers use two indices derived from the other citations of patents citing patent p: If patents that cite p also cite the older literature that p is referencing, p did not disrupt the technology. If however p starts a new literature that does not cite pre-existing patents anymore, p is classified as disruptive. I simplify their exact specification by not counting the citations between patent p, its cited patents and patents citing p. Instead, I define the citation year gap CYG by observing the average filing year of the other citations from patents citing p.

$$CYG_p = \frac{\sum^C \bar{t}_{o;c}}{C_p} - t_p \tag{1}$$

where p is the patent of interest, c indexes patents filed up to five years after patent p and citing patent p and o indexes the other patents cited by c. $\bar{t}_{o,c}$ is the average year of the patents o cited by citing patent c and t_p is the filing year of the original patent. Thus, the citing year gap CYG is the difference between the average year referenced by the patents cited by the patent citing p. A positive number means that patents citing p on average reference patents filed after the patent of interest p, i.e. they are referencing new research instead of the old patents that p is based on. Thus, this measure intuitively is very close to the original by Funk and Owen-Smith (2017). However, because it is essentially continuous, it outperforms the original in small samples, where the fact that most patents only have one or two citations really matters: Indices based on referenced patent count only have a couple of values in practice, making their movements quite jumpy. I prefer CYG_p as a measure since I follow "disruptiveness" within (sometimes) small technology classes, not the entire patent sample, unlike Funk and Owen-Smith (2017).

2.4 Aggregate Trends

To understand what drives the aggregate decline in disruptive research, I split the data into different time consistent technology classes provided by PATSTAT. I create a balanced panel of 75613 IPC classes from 1980 to 2010. Figure 2 reports the trend of CYG for all technology classes. Most IPC classes mirror the aggregate downward trend: Patents filed in 1980 get cited by patents that reference work on average roughly 1 year older than the original patent. In 2010, the number has increased to roughly 5. Though most IPC classes experience a decline in the CYG, the difference between the most disrupted and the most incremental technologies is rising, with some IPC classes even exhibiting rising disruptiveness as measured by the CYG during the 90s. A major factor in this is the revolution in ICT technology: Of the 25 IPC classes with the highest CYG in 2010, 19 are categorized as telecommunications and 9 among those as "transmission of digital information", 3 more are in "computing and image processing" and another 2 are in "games (including video)". The measure thus produces sensible results. Table 1 reports summary statistics for the IPC class panel.

2.5 Effects of Disruptive Innovations

In the baseline specification, I define a whole technology field as disrupted using the share of citations of its disruptive patents: If patents with a positive CYG gain 50% or more of all citations over the next 5 years, I mark the technology class as disrupted. The frequency of such events declines from 12% (patent weighted avg.) in 1980 to 3% in 2010. However, concurrent with the US technology and productivity boom at the end of the century (Fernald, 2015; Garcia-Macia, Hsieh and Klenow, 2019), the frequency of disruptions increases between 1992 and 1998. The advantage of such a definition of disruption – based on the share of disruptive patents in citations – is that it does not presuppose anything about the future of the technology field: Both declining and rising fields can in principle be disrupted.

To understand what happens after such a disruption event with the IPC class, I perform an event study: I match disrupted IPC classes to never-

Figure 2: *Aggregate Evolution of Disruptive Innovation*

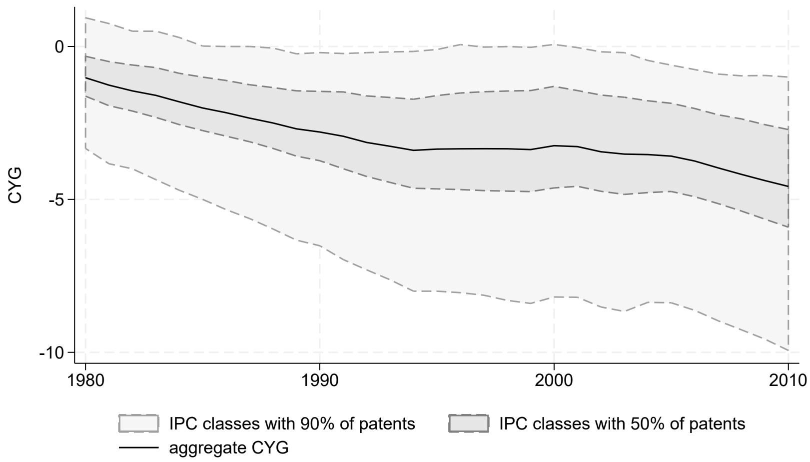

Notes: Average CYG per technology class over time. CYG measures how disrupted a technology is and is defined in (1) and discussed subsequently. In the aggregate over all IPC classes, the measure declines from -1 to almost -5. IPC classes worth 50% of all patents move in the dark gray area, IPC classes worth 90% move in the light gray area. The aggregate behavior is not driven by outlier IPC classes, declining CYG is widespread. However, some IPC classes increase their average disruptiveness, especially during the 90s. These are almost exclusively ICT-related.
Sources: PATSTAT (European Patent Office).

disrupted IPC classes in the same year and compare their evolution around the disruption event. Apart from exact matching on the year, I perform Mahanalobis distance matching on the CYG of the four years prior to the disruption, the number of new citations during the disruption year and the two years prior and the number of citations gained one year before the disruption by the inventor cohort that entered the economy 5 years prior to the disruption. Table 1 reports summary statistics for the two groups prior to matching as well as a comparison of the control and treatment group together with significance tests.

After the matching procedure, I obtain a sample of 1631 disrupted IPC classes and their nearest neighbor as a control. This is a substantial reduction from the 42283 IPC classes disrupted once and is mainly due to the difficulty

Table 1: *Summary Statistics on IPC classes before and after Matching*

	Panel 1: Before Matching			Panel 2: After Matching		
	Controls	Disrupted	Difference	Controls	Disrupted	Difference
CYG_{T-1}	-5.585	-3.441	2.144	-4.031	-3.917	0.114
	(4.231)	(3.821)	(0.044)	(2.996)	(3.205)	(0.109)
CYG_{T-2}	-5.485	-3.742	1.743	-3.907	-3.843	0.064
	(4.148)	(3.919)	(0.048)	(3.006)	(3.230)	(0.109)
CYG_{T-3}	-5.386	-4.008	1.378	-3.813	-3.783	0.029
	(4.067)	(3.903)	(0.052)	(3.048)	(3.266)	(0.111)
CYG_{T-4}	-5.278	-4.105	1.174	-3.752	-3.662	0.090
	(3.976)	(3.866)	(0.057)	(3.213)	(3.368)	(0.115)
$nr_{citations}(T)$	4.820	5.322	0.502	24.855	22.311	-2.544
	(65.112)	(8.486)	(0.317)	(25.414)	(27.623)	(0.929)
$nr_{citations}(T-1)$	4.820	3.186	-1.634	23.901	22.973	-0.928
	(65.112)	(7.374)	(0.317)	(22.709)	(23.086)	(0.802)
$nr_{citations}(T-2)$	4.391	2.544	-1.847	21.021	20.265	-0.755
	(59.560)	(6.494)	(0.290)	(20.732)	(20.128)	(0.716)
$cum.nr_{citations}^{cohortT-5}(T-1)$	1.187	0.999	-0.188	7.306	7.484	0.178
	(16.565)	(2.975)	(0.081)	(8.650)	(9.222)	(0.313)
Observations	1,477,476	42,283	1,519,759	1,631	1,631	3,262

Notes: Unit of observation: harmonized IPC class first disrupted in year T. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. CYG measures how disrupted a technology is and is defined in (1) and discussed subsequently. $nr_{citations}(T)$ denotes the number of citations to the technology class at time T. $cum.nr_{citations}^{cohortT-5}(T-1)$ refers to the total number of citations earned by inventors that entered 5 years before the disruption up to year $T-1$. In the population, disruptions happen in already less incremental IPC classes, measured by past CYG. Soon to be disrupted IPC classes have slightly less citations ex ante, because very small IPC classes are often tagged as disrupted with the definition used. After the matching procedure, the differences are controlled for. It is worth noting that matching mainly works for larger, well cited IPC classes and the matched sample reduces substantially.
Source: PATSTAT (European Patent Office).

of finding matches for small, less cited and already less incremental IPC classes. The difference in CYG between unmatched disrupted IPC classes and their potential control is large and does not allow to find matches for all disrupted IPC classes despite the large number of potential candidates.

The event study itself is estimated using OLS

$$y_{t^r;i} = \sum_{r=-5}^{r=15} \beta^{t^r} t_i^r + \Theta_i + u_{t^r;i} \tag{2}$$

where Θ_i is a matched pair fixed effect for pair i, t_i^r is relative time since the disruption and y stands for different outcome variables of interest. To capture the effect of disruptions on future disruption, I use CYG and the disruption dummy to capture the likelihood for consecutive disruptions. To study the effect of disruption on future innovation, I use the number of citations to patents in each field per year $nr_{citations}(t)$. To study the effect of disruptions on existing inventors, I follow the careers of the inventor cohort that entered 5 years before the disruption (minus the disrupting inventors). I track the citations that these inventors gain every year $nr_{citations}^{cohortT-5}(t)$.

Figure 3 reports the results of these estimations. Both the probability of a consecutive disruption (Panel 3a) and the CYG (Panel 3b) increase sharply with a disruption. Both trend downwards after the disruption, as does the CYG of undisrupted IPC classes. Disrupted IPC classes gain patent citations after a disruption and receive roughly $\frac{1}{3}$ more citations after 15 years (Panel 3c). Inventors who were not involved in disruptive innovation lose citations: After 15 years, the inventor cohort that entered 5 years before the disruption gets cited roughly 30% less every year than inventors in undisrupted IPC classes.

Especially the effect on existing inventors is affected by substantial measurement errors and should be treated as a lower bound estimate: The PATSTAT data does not come with inventor IDs (and neither do other patent data sets). Instead, it contains the names as written into the fields "inventor" and "applicant" on the patent. Thus, an important step whenever using PATSTAT inventors is to identify spelling mistakes and variants for which I improve upon Peeters et al. (2010); Magerman, Van Looy and Song (2006) and then disambiguate different inventors and firms with the same name, for which I build upon Toole, Jones and Madhavan (2021); Li et al. (2014). I detail these data treatments in Appendix A.

The results from this analysis present additional facts that the model should be able to explain:

6. Technologies' trend towards incrementalism is reversed by discrete, high impact, disruptive inventions.

7. Disruptive inventions increase the likelihood for consecutive disruption.

8. Disruptive inventions increase the citations earned by consecutive incremental research.

9. Existing inventors lose when others disrupt their field.

3 Model

This section develops a tractable endogenous growth model that captures the stylized facts discussed in Section 2. I adapt a standard dynamic equilibrium endogenous growth model, more specifically Akcigit and Kerr (2018), to add an inventor labor market and link inventors to technology clusters. Inventors and firms are also linked to the two types of innovation in the model: Incremental innovation adds product quality and disruptive innovation makes incremental inventors obsolete but increases the invention step size for future inventors. Firms cannot directly invest into either type of R&D, but instead have to hire inventors. When firms hire them, their innovation is directed towards these inventors' technologies and this direction is observable by all. Thus, firms interact strategically on the labor market for inventors – poaching inventors that pursue threatening research and building up their own inventor portfolio. The value function of firms pursuing incremental research illustrates the poaching incentive and thus the central innovation of the model:

Figure 3: *Effect of Disruption on IPC class*

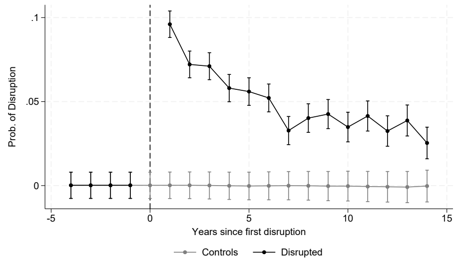

(a) Probability of new disruption

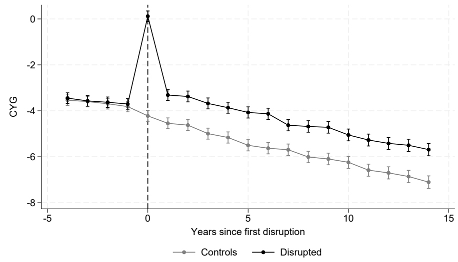

(b) Average CYG in the technology

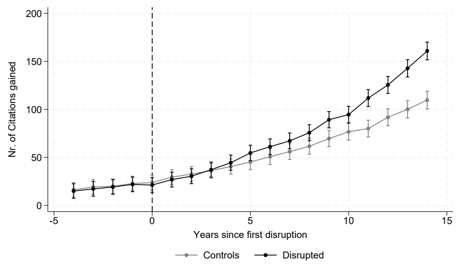

(c) Gained citations IPC class

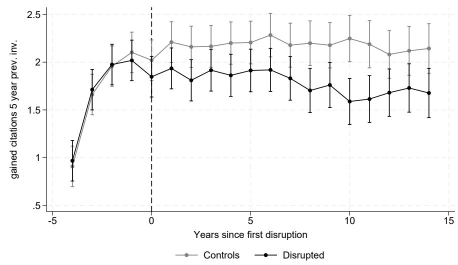

(d) Gained citations old inventors

Notes: Result of an event study in the nearest neighbor matched sample of IPC classes. Panel 3a depicts the evolution of the probability of a new disruption after the first disruption. Panel 3b depicts the evolution of the citation year gap as defined in section 2.3 as a measure of the overall "disruptiveness" of average patents in the technology class. Panel 3c tracks the number of citations in an IPC class per year. Panel 3d shows the citations gained by the inventors that had entered the IPC class 5 years before the disruption. *Sources*: PATSTAT (European Patent Office).

$$rV_f(\widetilde{q}_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis}) - \dot{V}_f(q_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis}) = \max_{\{\dot{\lambda}_f^{inc}; w_f^{dis}\}} \pi(q_f) + V^{patent}\lambda_f^{inc} - w(\lambda_f^{inc}) * \lambda_f^{inc}$$

$$-\Lambda^{dis} * \Delta V_f(q_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis}) - w_f^{dis} * \lambda_f^{dis}(w_f^{dis}) - c(\dot{\lambda}_f^{inc}) \quad (3)$$

Six terms define the value of an incremental firm: First, its patent portfolio, which generates profits $\pi(q_f)$. Second, the firm's current inventor portfolio creates new patents at rate λ_f^{inc}. In contrast to other endogenous growth models, this rate is not a choice variable, but depends on past successes in inventor recruitment. Third, these inventors demand wage $w(\lambda_f^{inc})$. Fourth, at rate Λ^{dis}, the current technology is disrupted and the firm loses a part of its value $\Delta V_f(q_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis})$. I will make the most conservative assumptions about this term and restrict losses to future inventions only, i.e. existing patents remain unaffected. Fifth, the firm might pay a wage to disruptive inventors w_f^{dis}, even though they do not produce anything for the firm: A high wage attracts disruptive inventors which decreases the rate of disruption Λ^{dis}. Sixth, the firm chooses $c(\dot{\lambda}_f^{inc})$, i.e. the costs it wants to incur to match with additional incremental inventors. Comparing eq. (3) to a more standard model, the firms cannot choose the rate of inventions, since the stock of inventors and their expected output is given at time t. Instead they choose their recruitment effort for incremental inventors (which determines $\dot{\lambda}_f^{inc}$) and the wage they pay to disruptive inventors who stop disrupting w_f^{dis}.

The next subsections will derive the parameters and functions that determine the firm's decision: the value of an incremental patent V^{patent} (section 3.1), the wage that incremental inventors can demand $w(\lambda_f^{inc})$, the relationship between poached disruptive inventors and the wages that firms post for them ($\lambda_f^{dis}(w_f^{dis})$) which also affects the rate of disruption in the technology field (Λ^{dis}) and the costs of hiring incremental inventors $c(\dot{\lambda}_f^{inc})$. Firm/inventor variables are lower case and technology field level aggregates are capitalized.

3.1 Demand and Value of Patents

Consumers are part of a representative household and derive logarithmic utility from consuming a final good (Y) in continuous time. This final good is the numeraire good.

$$U = \int_0^\infty e^{-rt} \ln(Y(t))dt$$

Consumers discount the future with factor r. They neither face a tradeoff between leisure and consumption, nor do they experience inequality. Households evenly share income from all sources between their members.

The final goods industry produces the consumption good using unskilled labor and a variety of intermediate inputs and sells it to consumers. The industry produces according to

$$Y(t) = \frac{1}{1-\beta} L_c^\beta(t) \int_0^1 q_j^\beta z_j^{1-\beta} dj \tag{4}$$

where q_j is the quality of good j, z_j is its quantity and $L_c(t)$ is the unskilled labor expended in final goods production. If all product qualities are fixed, the production function exhibits constant returns to scale in labor and intermediate inputs. When product quality q_j also increases, the production function exhibits increasing returns to scale. Each product j is produced by intermediate firms, who improve its quality via research and are the primary actors of the model.

The final goods industry is a price taker, consisting of a multitude of small competing firms. Hence, its inverse demand for any one intermediate good is $p_j = L_c^\beta(t) * q_j^\beta * z_j^{-\beta}$. A monopolist firm producing the intermediate good j with production function $z_j = \bar{q}l_j$ (where $\bar{q}$ is the average quality in the economy) would generate profits of

$$\pi(q_j) = \pi^*_{mon} * q_j = [L_c(t) * (1-\beta) * \beta^\beta(1-\beta)^{1-2\beta}] * q_j \tag{5}$$

Thus, a monopolist's profits are a linear function of quality and (from the viewpoint of the firm) an exogenous factor called π throughout the rest of the paper.

Each product corresponds to one technology field (like "telecommunications" or "electricity generation"). There are multiple firms in each technology field, which cannot do research themselves and have to hire inventors to increase product quality. Within each of these fields, there is a hierarchy of technology clusters, which denote better and different ways of producing the product in question. E.g. the clusters "telegraphy" or "satellite communication" in the field "telecommunications". These clusters are areas of expertise for individual inventors, who cannot be experts in whole fields or even all sciences. Within each field, there are old, obsolete clusters ("telegraphy"), a currently active cluster ("satellite communication") and as of yet still unknown future clusters. Firms are linked to a specific technology cluster through the inventors they have hired. Inventions in higher clusters have a bigger impact on product quality.

I depart from Akcigit and Kerr (2018) in how competition between firms with the same product (i.e. within the same technology field) works: I assume that patented inventions are additive in quality, that is each represents an independent increase in quality of ω^c. The quality of the produced good is the sum of the quality improvements of the patents a firm either holds or licenses. Production is still represented by a two stage game, during the first stage of which firms license their patents to each other. In the second stage there is only one actual producer, who licenses all existing patents and charges the monopoly price. This variation on the two stage game makes finished inventions themselves safe: Firms' patents will retain value even after another firm has made an additional invention. This assumption simplifies the value function of firms in light of the two new assets I introduce (incremental inventors and disruptive inventors).

With these assumptions, firm value $V_f(\widetilde{q}_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis})$ is additively separable into the value of the firm's patent portfolio (which creates passive rents, but does not affect any decision of the firm and is not at risk from disruption) and the firm's inventor portfolio, which is affected by the firm's decisions

$$V_f(\widetilde{q}_f, \lambda_f^{inc}, \lambda_f^{dis}, \Lambda^{dis}) = V^{patent} * \widetilde{q}_f + V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis})\lambda_f^{inc} \tag{6}$$

where $V^{patent} * \widetilde{q}_f$ denotes the value of the patents of the firm and $V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis})$ is the value of an inventor portfolio with a patent arrival rate of 1.

With these additions, the value of a patent is

$$V^{Patent}(c) = \omega^c * \frac{\pi}{r} \tag{7}$$

where ω^c is the quality increase of an invention, dependent on its technology cluster c. ω describes the growth in step size generated by disruptive inventions and thus is the main driver of aggregate growth. Eq. 7 captures stylized fact (8): Disruptions increase the value of future incremental research (by factor ω).

Figure 4: *Innovation Example*

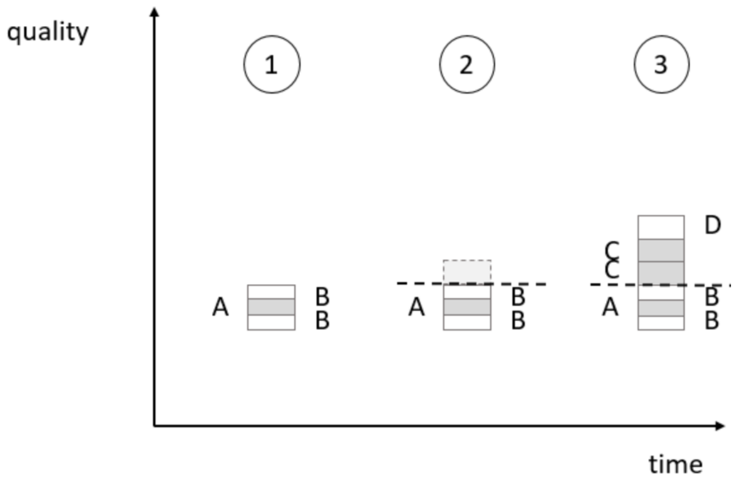

Notes: An example of the evolution of a technology field before (1), during (2) and after (3) a disruptive invention. Firms A and B pursue incremental innovation with a constant step size before the disruption. As a result, productivity growth declines as ideas are getting harder to find. After the disruption, A and B's inventors can no longer contribute to the new cluster. New firms C and D start innovating with increased step size.

This technology setup replicates key features of Cowen and Southwood (2019) and Olmstead-Rumsey (2019) insofar as productivity growth is slowing down within specific technology clusters (or tasks) through the channel of patent quality. However, for modeling convenience, patent quality is stagnant in undisrupted technology fields and increasing in technology fields with regular disruption.

3.2 Incremental Inventors

Both incremental and disruptive inventors enter any technology field at an exogenous rate H^{inc} or H^{dis}. Inventors have a quality x_i drawn randomly from the uniform distribution between 0 and 1. The incremental inventor labor market connects entering incremental inventors with the set of existing firms. There exists a mass 1 of incremental firms which draw a research quality y_f, also from the uniform distribution between 0 and 1. After a firm and an inventor meet, their types are revealed and they Nash bargain over the inventor's wage. The pair produces patents at rate $y_f x_i$ if they match. This supermodular production function means firms and inventors depend on the quality of their counterpart. At any point, there is a risk δ that the inventor exits the economy.

Wages in this market are determined by Nash bargaining within matches. Since both can freely terminate and renegotiate the contract, they bargain over the wage of the current period only. However, neither side has a credible outside option. The inventor will not get additional matches and will thus be stuck negotiating with the firm again. The firm's matches are independent from each other since the specific vacancy is destroyed after matching whether the firm accepts or not. The match produces $\pi\omega^c y_f x_i$ at any time, where π is the constant profits that can be expected from increasing productivity, ω^c is the productivity improvement of one incremental invention which depends on the index of the current technology cluster c and $y_f x_i$ is the rate at which new inventions are generated, a function of firm quality and inventor skill. Since both outside options are 0, the entire output of the match is the surplus, which will be divided between firms and inventors according to the bargaining parameter α. This yields the wage bill for incremental inventors for firm f as:

$$w^{inc}(\lambda_i) * \lambda_i = (1-\alpha)\pi\omega^c * \lambda_i = (1-\alpha)\pi\omega^c * y_f \int x_i \,\mathrm{d}x_i \tag{8}$$

In equilibrium, incremental firms and inventors will always work together to produce patents, with the inventor i earning share $(1-\alpha)$ of the expected profits as wage and the firm f earning the remaining share of α.

To actually match with incremental inventors, firms create vacancies ξ_f at cost $y_f * c_\xi$. Firms of lower research quality thus have lower costs, which keeps them in the market. Other setups are of course possible, but in the interest of an analytic solution, I stick to the most simple version of the inventor labor market. After firms have created their vacancies, entering inventors match with a random vacancy. Inventors do not match again, even if they reject this match and unmatched vacancies are destroyed. Appendix B discusses the effect of these simplifications. Given these assumptions, any specific firm matches with incremental inventors according to $\eta_f = \xi_f * \rho_f = \xi_f * \frac{H^{inc}}{\Xi^{inc}}$. Using the separability of the value function and the linearity of non-patent derived value with respect to λ_f^{inc} (eq. 6), maximizing with respect to $\dot{\lambda}_f$ yields

$$\frac{\partial c(\dot{\lambda}_f^{inc})}{\partial \dot{\lambda}_f} = y_f c_\xi * \frac{1}{\frac{1}{2} * y_f} \frac{H^{inc}}{\Xi^{inc}} \stackrel{!}{=} V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis}) \tag{9}$$

where $y_f c_\xi$ is the cost of creating a vacancy and $\frac{2}{y_f} \frac{H^{inc}}{\Xi^{inc}}$ is the number of vacancies needed to attract inventors with a total patent arrival rate of 1: The expected skill of attracted inventors for any firm is $\frac{1}{2}$ and is then multiplied with the firm's research quality y_f to yield the expected arrival rate of patents. $\frac{H^{inc}}{\Xi^{inc}}$ is the ratio of entering inventors to vacancies and thus the success rate of any individual vacancy. $\frac{H^{inc}}{\Xi^{inc}}$ thus captures the congestion externality usually found in search and matching labor markets: By increasing the number of vacancies, each firm increases the overall competition for inventors and reduces the efficacy of each individual vacancy. Firms will increase their vacancy creation until the (aggregate) success rate of any individual vacancy is so low that eq. (9) holds. Note that eq. (9) implies that there is no value in being a "new" incremental firm, i.e. one without researchers: Incremental firms compete with each other for inventors by increasing the number of vacancies so much that all gains from additional inventors are expended in vacancy creation.

As a result of the cost function for vacancies, each firm creates the same amount of vacancies and receives the same amount of inventors $\eta_f^{inc} = \frac{H^{inc}}{1}$, while all firms' inventors exit at rate δ. The total skill of all accumulated inventors is

$$x_f = \frac{1}{2}\frac{H^{inc}}{\delta}(1 - e^{-\delta t}) = X^{inc} \tag{10}$$

where x_f denotes the total skill of the inventors of a firm and X^{inc} those in the technology field (identical since there is a mass 1 of firms of identical size).

The assumptions made here follow the standard search and matching labor market setup as surveyed e.g. in Rogerson, Shimer and Wright (2005), with the possible exception of the simplified matching function, which serves to keep the value functions analytically solveable.

3.3 Disruption and Disruptive Inventors

Like incremental inventors, disruptive inventors enter each field at rate H^{dis} and draw a skill level x_i from the uniform distribution between 0 and 1. They also immediately draw a random incremental firm with which they are now matched. However, they do not produce anything together. Instead, each disruptive inventor produces a disruptive invention with rate x_i as long as his matched incremental firm does not offer a high enough wage to poach him from this activity.

Disruptive inventions are meant to represent prototypes of future production technologies. Whenever disruptive inventors are successful, a new technology cluster is born and the old cluster becomes obsolete. Old incremental inventors can no longer contribute to product quality after a disruption, but disruptive inventors can immediately work on disrupting the new technology again. To represent the first mover advantage, the disrupting inventor also earns γ patents in the new cluster, which represent the private gains from disruptive innovation. The higher this parameter, the higher the incentive to keep making disruptive inventions.

Incremental firms view the rate of disruption as a risk to their stock of incremental inventors. To avoid modeling a fully fledged secondary labor market where disruptive inventors repeatedly bargain with incremental firms, I assume that disruptive inventors hold all bargaining power and can extract the full match surplus from the incremental firms they are negotiating with. Effectively, they can make take-it-or-leave-it offers. An incremental firm's maximum willingness to pay off a disruptive inventor i is given by the value of the firm's portfolio of incremental inventors and has to exceed the value of potential disruptive innovation, i.e.

$$V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis}) * X^{inc} * y_f * x_i \geq \gamma \omega V^{Patent}(c) * x_i \tag{11}$$

where $V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis}) * X^{inc} * y_f$ denotes the value of the incremental inventors the firm holds ($\lambda_f^{inc} = X * y_f$) and x_i is the reduction in the rate of disruption that the inventor represents (his patent arrival rate x_i). The right term represents the value of working on disruptive innovation, which is more valuable than a current incremental patent by the factor of $\gamma\omega$ (the first mover advantage and the step size gain from a disruptive innovation). Solving for y_f gives the cutoff research quality below which firms will not poach disruptive inventors. Over time, the number of incremental researchers per firm (X^{inc}) and thus the cutoff value y_f^* falls. More and more firms will poach their matched disruptive inventors. Additionally, the value of incremental inventors is also rising as the risk of disruption becomes smaller, acerbating the process.

The assumption that disruptive inventors gain all the surplus from bargaining is less restrictive than it originally seems: Rogerson, Shimer and Wright (2005) discusses why in matching labor markets, matching behavior does not depend on the surplus negotiation. Thus, no matter the assumption about the bargaining parameter, whether or not disruptive inventors accept offers from incremental firms and vanish from the market depends on the above inequality. The assumption is only relevant for the value functions of incremental firms, which are kept tractable this way.

3.4 Main Theoretical Results – Equilibrium within Technology Field

The value maximization of incremental firms drives the evolution of individual technology fields. Due to the assumptions made about demand and the technology process in section 3.1, the income from patents has no effect on the decision variables of the firm and firms actually maximize the value of their inventor portfolio $V_f^{inv}(1, \lambda_f^{dis}, \Lambda^{dis}) * \lambda_f^{inc}$. Sections 3.2 and 3.3 described the labor markets for incremental and disruptive inventors that these firms navigate. Inserting equations (7) to (11) into the value function of the incremental firm (eq. 3) yields the value of a firm's inventor portfolio as a function of the aggregate stock of incremental inventors X. Firms with a research quality below the cutoff y_f^* do not participate in the labor market for disruptive inventors. As such, an analytical solution for their value function is not necessary. I instead derive the value of firms at or above the cutoff to understand the behavior and prevalence of poaching firms:

$$\begin{aligned} rV_f^{inv}(1, \lambda_f^{dis}, X^{inc}) = \quad & \underbrace{\frac{\pi}{r}\omega^c * \alpha}_{\text{new patents net of inv. wages}} \quad - \underbrace{\delta V_f^{inv}(1, \lambda_f^{dis}, X^{inc})}_{\text{inv. exit}} \\ & - \underbrace{\Lambda_{max}^{dis} \frac{\gamma\omega\pi * V_f^{inv}(1, \lambda_f^{dis}, X^{inc})}{V_f^{inv}(1, \lambda_f^{dis}, X) * X^{inc}}}_{\text{disruption risk}} \\ - \underbrace{\lambda_f^{dis} V_f^{inv}(1, \lambda_f^{dis}, X^{inc})}_{\text{wages to poached inv.}} & + \underbrace{\frac{\partial V_f^{inv}(1, \lambda_f^{dis}, X^{inc})}{\partial X^{inc}}(H^{inc} - \delta X^{inc})}_{\text{increase in poaching by others}} \end{aligned} \tag{12}$$

where the first term denotes the value from the invented patents minus the wages paid to incremental inventors and the second term captures the losses from incremental inventors exiting at rate δ. The risk of disruptive inventions declines with the share of non-poaching firms $y_f^* = \frac{\gamma\omega\pi}{V_f^{inv}(1,\lambda_f^{dis},X^{inc})*X^{inc}}$ as described in section 3.3. Since the firm is itself above the cutoff, it will have to pay wages of $\lambda_f^{dis}V_f^{inv}(1,\lambda_f^{dis},X^{inc})$ to the disruptive inventors it has itself poached. The last term denotes the value change as all firms aquire more inventors and the share of non-poaching firms declines.

Solving eq. (12) yields the value of poaching firms and their resulting strategy. Since non-poaching firms by definition do not interfere with disruptive innovation and all firms have the same hiring behavior on the market for incremental inventors, this characterizes the equilibrium within a technology field as

$$V_f^{inv}(1,\lambda_f^{dis},X^{inc}) = \frac{\pi}{\delta+r+\lambda_f^{dis}}[\alpha - \frac{2\gamma\omega}{X^{inc}}{}_2F_1\left[\begin{matrix}1 & \frac{r+\lambda^{dis}}{\delta}+1 \\ & \frac{r+\lambda^{dis}}{\delta}+2\end{matrix};1-\frac{X^{inc}}{X^{inc}_{max}}\right]] \tag{13}$$

where ${}_2F_1\left[\begin{matrix}a\ b\\ c\end{matrix};z\right]$ is the Gaussian hypergeometric function.

$$\Lambda^{inc}(X^{inc}) = X^{inc} * \int_0^1 y_f \,\mathrm{d}y_f = \frac{1}{2}\frac{H^{inc}}{\delta}(1-e^{-\delta*t})\frac{1}{2} \tag{14}$$

$$\Lambda^{dis}(X^{inc}) = \int_0^{y_f^*}\frac{1}{2}\frac{H^{dis}}{\delta}\,\mathrm{d}y_f = \frac{1}{2}\frac{H^{dis}}{\delta}y_f^* = \Lambda^{dis}_{max} * \frac{\gamma\omega\pi}{V_f^{inv}(1,\lambda_f^{dis},X^{inc})*X^{inc}} \tag{15}$$

The value function of an incremental inventor has two components: The value of an inventor if the rate of disruption were 0 ($\frac{\alpha\pi}{\delta+r+\lambda^{dis}}$) and a share capturing the changing risk of disruption. An incremental inventor's value is lower if $\gamma*\omega$ rises, i.e. if there is a larger first mover advantage for the disrupting inventor or if disruptive inventions create larger step size increases. Higher returns to disruptive innovation mean that fewer incremental firms can pay sufficient wages to poach disruptive inventions. The faster the stock of incremental inventors grows after a disruption, the faster the value of inventors increases, too. The growth of this stock (and thus also Λ^{inc}) is governed by the rate at which new inventors enter and the rate at which inventors exit the economy again. Directly after the disruptive innovation, at $t = 0$, all incremental inventors have been made obsolete and none are active. The probability of a consecutive disruption Λ^{dis} moves in the opposite direction and declines over time as the number of incremental inventors and their value increases.

Figure 5 describes the equilibrium within one technology sector as a function of the time since the last disruption. Over time, the number of incremental inventors increases, the share of non-poaching firms decreases and the chance of disruption declines while the output of incremental innovation increases.

3.5 Aggregate Growth and Steady State

The different technology fields and patents are linked by aggregate demand as in section 3.1 and the labor market for unskilled labor. Apart from the intermediate monopolists, unskilled labor is also demanded by the final goods sector, which will optimize their labor and intermediate goods intake and through this set the wage rate. Optimizing eq. (4) with respect to labor and inserting the equilibrium on the intermediate goods market (eq. 5) gives the optimal wage as

$$w = \beta^{\beta}(1-\beta)^{1-2\beta} * \bar{q} \tag{16}$$

i.e. the final goods industry will adjust its labor demand to achieve a wage rate as a multiple of the average quality $\bar{q}$ in the economy. The precise multiple is dictated by labor's output elasticity β.

Aggregate growth is driven by technological progress via the two different types of innovation: Incremental inventions improve the average quality of the products in the economy and ultimatively increase the utility of consumers. Disruptive progress increases the value of future incremental progress and ensures long-term growth: Without disruptive innovation, the economy still grows as new incremental inventions increase quality, but growth as a percentage of GDP declines because incremental inventions can only create linear growth.

On the steady state growth path, the number of technology fields with

Figure 5: *Equilibrium within a Technology Sector*

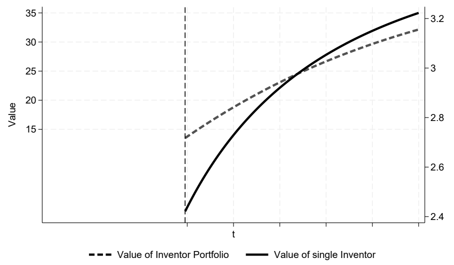

(a) Value of Incremental Inventors

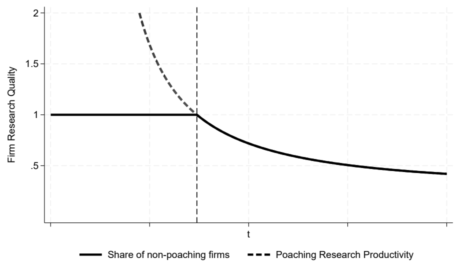

(b) Share of Non-Poaching Firms

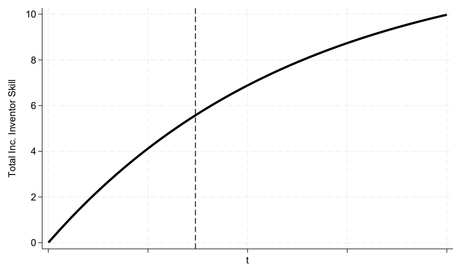

(c) Total Skill of Incremental Inventors

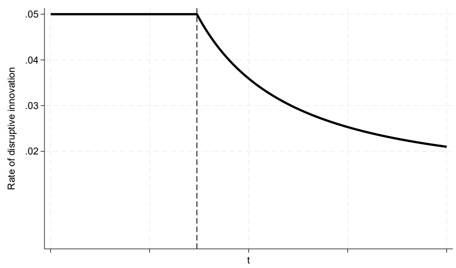

(d) Rate of Disruption

Notes: Description of the equilibrium path within a technology field. Panel 5a shows the evolution of the value of a single incremental inventor with $\lambda_i = 1$ and of an entire portfolio of such inventors for a poaching firm. Panel 5b shows the resulting cutoff firm quality above which firms will start to poach disruptive inventors: It is declining over time as the inventor portfolios get larger and (once the first firms poach) individual incremental inventors get more valuable. Panel 5c and Panel 5d describe the aggregate incremental inventor skill and the rate of disruptive inventions.

any specific rate of disruption $N_{field}(\Lambda^{dis})$ is stable, which keeps both the aggregate rate of disruptions and the rate of incremental inventions constant. Note that Λ^{dis}_{field} fully characterizes a field, since it is monotonely dependent on t (time since last disruption), as is Λ^{inc}_{field}. Λ^{dis}_{field} or t both define the type of a field.

In the steady state, the inflows into any type must equal the outflows. Fields that are disrupted move to $\Lambda^{dis}_{field}(t=0)$ and t_{field} increases linearly for all other fields. This translates to the differential equation $-N_{field}(t_{field}) * \Lambda^{dis}_{field}(t_{field}) = \dot{N}_{field}(t_{field})$. Figure 6 compares the actual age distribution in 1990 and the distribution after 20 years of simulation, which is extremely close to the steady state.

4 Main Results – Explanatory Power of the Model

Since aggregate growth is a function of disruptive inventions, $\Lambda^{dis}_{field}(t)$ determines aggregate growth, where t is "age" of the technology field, i.e. the time since the technology field was last disrupted. The model laid out above argues that aggregate growth can decline as a result of technology fields aging in this way, since incremental firms grow large and attempt to hinder disruptive innovation. To gauge the size of this effect, I calibrate the model to the global patenting landscape in 1990 and simulate the next 20 years. This yields the predicted evolution of $\Lambda^{dis}_{field}(t)$ without any changes in the underlying parameters.

Key parameters of the model can be observed in patenting data: The exit rate of non-single-patent inventors δ is 0.15 in 1990. To set the average size of a technology field $X^{inc}_{m}ax$, I take the maximum number of patents per year in each technology field and take the average. I measure the average number of firms in a technology field as the firms with a patent both before and after 1990 (this mass of firms is set to 1 in the theoretical model). I take Λ^{dis}_{max} from the event study performed in section 2.5. I set α to the labor share in the overall economy, since I have no way to measure inventor wages. I set the real interest rate to 8%.

To gauge the age distribution t of the technology landscape in 1990, I use the previous 10 years of my data (1980-1990) and the disruption measure defined and discussed in section 2. Ages above 10 cannot be measured reliably, since the data is only complete since 1980. However, the additional effects of age above 10 are so small this does not affect the simulation.

Last, I estimate the only unconstrained parameter $\gamma * \omega$ to fit the model to the actual rate of disruption in 1990. Since only the $\gamma * \omega$ affect the evolution of the economy, it is impossible to disentangle the first mover advantage from the step size of the economy.

Figure 6 reports the results of the simulation and compares it to the actual data. Panel 6a reports the age distribution in 1990, 2010 and the result of the simulation in 2010. As predicted by the simulation, the age distribution shifts substantially towards higher ages, especially 10+ years since the last disruption. However, the change in the age profile is even more pronounced in the data than in the simulation. A similar picture emerges from Panel 6b, which shows the evolution of the share of IPC classes with ages 1-5 and 6-10. Both shares are declining substantially, however, there is a noted increase between 1992 and 1998, which coincides with the start of the tech and productivity boom in the US (Fernald, 2015; Garcia-Macia, Hsieh and Klenow, 2019). Figure 6d reports the result of the final exercise: Taking the parameters summarized in Table 6c and calibrating $\gamma * \omega$ to the disruption rate in 1990 is all that is necessary to simulate the further evolution of the economy. Comparing the simulated evolution of the rate of disruption with the actual evolution, the simulation can explain 52% of the decline from 1990 to 2010. The actual patent data shows more disruption than expected during the 90s and less during the 2000s.

The simulation can explain the majority of the observed decline without any parameter changes. This is not a refutation of the hypothesis that "ideas are getting harder to find", but it offers an alternative explanation. The stylized facts reported in section 2 also support the idea that firms and inventors face substantial technology risks from other firms' disruptions and that disruptive innovations lead to higher productivity research. Finally, the

Figure 6: *Counterfactual Simulation*

(a) Age Distribution

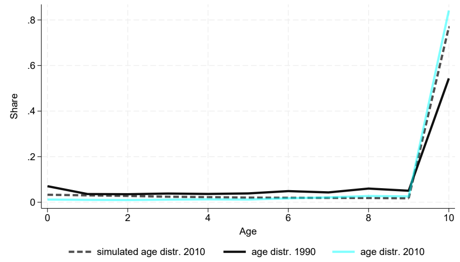

(b) Declining Share of Young Technologies

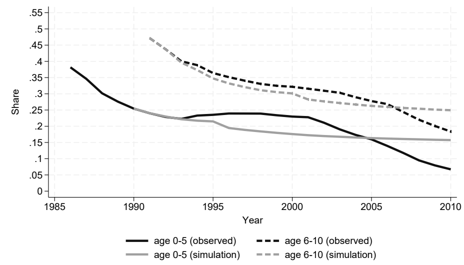

(c) Simulation Parameters 1990

Parameter	**Value**	**Source**
δ	0.15	observed
X_{max}^{inc}	28.5	observed
M_f	27.8	observed
r	0.08	
α	0.33	
Λ_{max}^{dis}	0.1	section 2.5
$\omega * \gamma$	7.83	**estimated**

(d) Observed and Simulated Disruptions

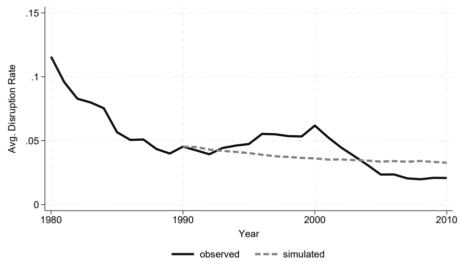

Notes: Panel 6a reports the distribution of technology ages in 1990 and 2010, as well as the 2010 age distribution resulting from the counterfactual simulation. Panel 6b shows the declining share of young technologies: While 25% of technologies were younger than 5 years in 1990, only 6% were so in 2010. Panel 6d compares the observed rate with the expected evolution of disruptions, given the parameters estimated for 1990 (Panel 6c).
Sources: PATSTAT (European Patent Office).

model as presented can explain the trend towards incrementalism without resorting to the assumption that it is exogenously given. There exists no counterfactual world technology frontier with which to directly rule out one of the explanations. However, it is consistent with the evidence that successful firms inhibit disruptive innovation and pursue incremental strategies. However, firms' behavior can be affected with policy, in contrast to exogenous technology trends. Section 5 discusses the impact of potential policy interventions.

5 Policy Implications

The economy presented in the baseline specification has several major decision points, only some of which the market economy handles efficiently. First, there is the demand of the final goods sector for intermediate products to turn into the final consumer product. The economy has a fixed number of products defined by how many technology fields there are and all of them are produced in equilibrium. However, the quantity produced is smaller than in the optimum because of the monopoly power of firms. This inefficiency depresses output by a fixed share, but has no impact on equilibrium growth rates. Second, firms have to hire incremental inventors to improve their product. Producers hire all incremental inventors. So, there is no inefficiency in this dimension. Third, disruptive inventors work on disrupting the economy and get poached by producing firms to prevent this. These poaching firms only have to outbid the private benefits of disruptive innovation. These do not include the productivity gains of future inventors. Currently existing incremental firms bear all costs from disruption and receive none of the benefits, thus they have a strong incentive to prevent disruption. A social planner that maximizes the utility of representative households makes a very different calculation: He weighs the value of getting inventors empowered by the disruptive invention in the future against the costs of losing all current inventors. A social planner might still arrive at the same conclusion as the market economy if consumers are sufficiently impatient: Empowering future inventors takes longer to pay off than current incremental inventions.

This highlights an important point about the tradeoffs involved in the decision about which type of research to pursue: Increasing long-run economic growth in this model requires unambiguously hurting the current generation. The currently living incremental inventors and firms have a vested interest in slowing economic growth. Fast productivity growth through disruption does not benefit them, but the inventors and firms who will enter the newly created cluster. This cannot always be solved via transfers: The current stock of incremental inventors is made obsolete, temporarily decreasing GDP growth. While it will eventually be rebuilt and growth will increase, many incremental inventors and firm owners that were hurt by the disruption will already have left the economy. Effectively, the current generations prefer to increase the level of economic activity through incremental inventions at the cost of economic growth. Of course, the linear technological progress of incremental improvements is still progress, but it means that the growth rate of the economy will decline.

The arrival rate of disruptive inventions in the economy is determined by the value of the stock of incremental inventors (eq. 11) vs. the value of a disruptive invention. Figure 7 reports the alternate equilibrium paths from changing select parameters in 1990. Policies that affect growth thus have to target four sets of parameters:

γ gives the first mover advantage of the disruptive inventor. This parameter captures how much of the disruptive invention the original inventing firm can appropriate (measured in "free" patents). Increasing this value increases the incentives for disruptive inventions and makes poaching more expensive. Many countries support disruptive (and incremental) inventors in incubators etc. to ensure that as many of these potential firms are as successful as possible, both through cash injection but also through transferring business knowledge and similar methods. If these activities increase the first mover advantage or business success rate of disruptive inventors, they benefit long run growth in the model.

r, δ, H^{inc} and y^{max} determine the value of the incremental inventor portfolio: A high r means firms discount the future profits from incremental inventors less. Inventor exit rate δ and incremental inventor entry rate H^{inc} determine how many incremental inventors exist and thus how valuable the firms' incremental inventor portfolio is going to be. y^{max} is the highest firm research quality in the economy and 1 by assumption in the above model. The higher the maximum firm research quality, the higher the incentive for these firms to poach disruptive inventors. Yet, high firm quality benefits incremental innovation, so lowering this parameter incurs clear costs.

With the above assumptions, the share of poaching firms $(1 - y_f^*)$ solely determines the share of poached inventors. However, a better working labor market for disruptive inventors could allow firms to poach more than "their" share of disruptive inventors and would cause aggregate growth to decline. Likewise, if the labor market was worse and incremental firms could not match with disruptive inventors at all, aggregate productivity would rise. There is an active literature on the question of whether startup acquisitions are welfare-enhancing (Cabral, 2018; Piazza and Zheng, 2019). My paper offers an additional argument for prohibiting such acquisitions.

ω captures the gains from disruptive innovation and is conceptually mostly a technology parameter. There are no downsides to having as high of an ω as possible: It increases the rate of disruptive innovations and also the gains from disruptions. However, it is unclear which policies can affect this parameter (and there is no reason why these policies should not already be implemented).

Figure 7: *Effect of Parameter Changes on Disruptive Innovation*

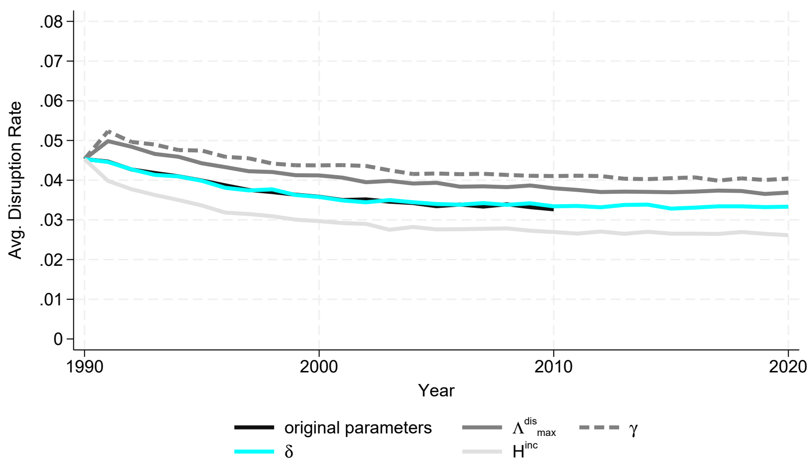

Notes: Effect of increasing various parameters by 10%. The entry rate of incremental inventors H^{inc} increases the value of patent portfolios and thus increases poaching. Both the stock of disruptive inventors Λ^{dis}_{max} and the first mover advantage γ increase the equilibrium path of disruptive innovation. The inventor exit rate δ decreases the equilibrium stock of incremental inventors but increases the speed with which the inventor stock reaches that equilibrium and has a negligible overall effect.
Sources: Own simulations.

6 Conclusion

The main contribution of the paper is to build an endogenous growth model around the difference between incremental and radical/disruptive innovation in which firms prevent disruption by poaching inventors. This mechanism jointly reproduces the decline in aggregate growth, the rise in incremental research (Park, Leahey and Funk, 2023; Funk and Owen-Smith, 2017; Kalyani, 2024) and the decrease in research productivity for specific tasks (Cowen and Southwood, 2019; Bloom et al., 2020). To achieve these results, I insert an analytically tractable search and matching labor market for inventors into an endogenous growth model and avoid the numerical solutions typically associated with these labor markets (Rogerson, Shimer and Wright, 2005; Hagedorn, Law and Manovskii, 2017). Firms hire specialist inventors to incrementally improve their technologies. Disruptive inventors threaten to make these technologies obsolete and devalue these inventor portfolios. Firms thus poach disruptive inventors to protect their investment. As they grow, incremental firms prevent an increasing share of technology disruption and become even more valuable, which raises the costs of disruption further. In aggregate, the economy stagnates due to a lack of technology disruptions.

To confirm the empirical relevance of the model's main mechanism, I perform an event study around disruptions in specific technology fields using PATSTAT from 1980 to 2010. I show that disruptive innovation increases subsequent patent citations in the field, but citations of already established inventors decline (both by roughly 30%). Disruptive inventions also increase the likelihood of subsequent disruptions, though the effect is decaying over time. The model also predicts such a pattern: Disruption destroys incremental firms and decreases poaching – at least until new incremental firms rise. Technology fields without disruptive inventions continuously become less relevant and less likely to be disrupted in the future.

To gauge the importance of the model's mechanism, I calibrate the model to the 1990 patent data and simulate the next 20 years of technology evolution. The simulation explains 52% of the observed decline in the rate of disruptive innovation and produces similar technology field age profiles. Between 1992 and 1998, the patent data shows an acceleration of disruptive innovation, which coincides with the IT boom and its resulting productivity and creative destruction in the US (Fernald, 2015; Garcia-Macia, Hsieh and Klenow, 2019). In contrast, the simulation expects a continuous decline, which describes the data after 1998 well.

The model implies several levers for policy. Interventions can make it more costly for incrementally innovating firms to poach by increasing the expected value of searching for disruptive inventions: Innovation prizes, support for innovative startups (through incubators etc.) or increasing the base number of disruptive inventors through education all fall under this category. Interventions can also target the poaching of disruptive inventors directly by e.g. preventing established firms from buying startups. Both measures will increase long-term growth. In my model, poaching or acquisitions change what type of research is conducted, which cannot be counteracted by a potential increase in research activity through acquisitions as controversially discussed in Cabral (2018); Piazza and Zheng (2019); Naidu, Posner and Weyl (2018).

7 Bibliography

Acemoglu, Daron. 2023. "Distorted Innovation: Does the Market Get the Direction of Technology Right?" *AEA Papers and Proceedings*, 113: 1–28.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. "Transition to Clean Technology." *Journal of Political Economy*, 124(1): 52–104.

Akcigit, Ufuk, and Nathan Goldschlag. 2023. "Where Have All the "Creative Talents" Gone? Employment Dynamics of US Inventors."

Akcigit, Ufuk, and Sina T. Ates. 2023. "What Happened to US Business Dynamism?" *Journal of Political Economy*, 131(8): 2059–2124.

Akcigit, Ufuk, and William R. Kerr. 2018. "Growth through Heterogeneous Innovations." *Journal of Political Economy*, 126(4): 1374–1443.

Andrews, D., C. Criscuolo, and Peter Gal. 2016. "The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy." *OECD Productivity Working Papers*, 5.

Antolin-Diaz, Juan, Thomas Drechsel, and Ivan Petrella. 2017. "Tracking the Slowdown in Long-Run GDP Growth." *The Review of Economics and Statistics*, 99(2): 343–356.

Arora, Ashish, Sharon Belenzon, Andrea Patacconi, and Jungkyu Suh. 2020. "The Changing Structure of American Innovation." *Innovation Policy and the Economy*, 20: 39–93.

Barkai, Simcha. 2020. "Declining Labor and Capital Shares." *The Journal of Finance*, 75(5): 2421–2463.

Bessen, James, and Eric Maskin. 2009. "Sequential innovation, patents, and imitation." *The RAND Journal of Economics*, 40(4): 611–635.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review*, 110(4): 1104–1144.

Breitmoser, Yves, Jonathan H. W. Tan, and Daniel John Zizzo. 2010. "Understanding perpetual R&D races." *Econ Theory*, 44(3): 445–467.

Bryan, Kevin A., and Jorge Lemus. 2017. "The direction of innovation." *Journal of Economic Theory*, 172: 247–272.

Cabral, Luis M. B. 2018. "Standing on the Shoulders of Dwarfs: Dominant Firms and Innovation Incentives."

Cowen, Tyler, and Ben Southwood. 2019. "Is the Rate of Scientific Progress Slowing Down?"

De Loecker, Jan, and Jan Eeckhout. 2017. "The Rise of Market Power and the Macroeconomic Implications." National Bureau of Economic Research Working Paper 23687.

Denicolò, Vincenzo. 2000. “Two-Stage Patent Races and Patent Policy.” *The RAND Journal of Economics*, 31(3): 488–501. Publisher: [RAND Corporation, Wiley].

De Ridder, Maarten. 2024. “Market Power and Innovation in the Intangible Economy.” *American Economic Review*, 114(1): 199–251.

Fernald, John. 2015. “Data for: Productivity and Potential Output before, during, and after the Great Recession.” *NBER Macro. Annual*, 29/1: 1–51.

Funk, Russell J., and Jason Owen-Smith. 2017. “A Dynamic Network Measure of Technological Change.” *Management Science*, 63(3): 791–817.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow. 2019. “How Destructive Is Innovation?” *Econometrica*, 87(5): 1507–1541.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War.* Princeton University Press.

Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii. 2017. “Identifying Equilibrium Models of Labor Market Sorting.” *Econometrica*, 85(1): 29–65.

Hall, Bronwyn H., and Rosemarie Ham Ziedonis. 2001. “The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995.” *The RAND Journal of Economics*, 32(1): 101–128.

Hopenhayn, Hugo, Gerard Llobet, and Matthew Mitchell. 2006. “Rewarding Sequential Innovators: Prizes, Patents, and Buyouts.” *Journal of Political Economy*, 114(6): 1041–1068.

Jaffe, Adam B. 2000. “The U.S. patent system in transition: policy innovation and the innovation process.” *Research Policy*, 29(4): 531–557.

Kalyani, Aakash. 2024. “The Creativity Decline: Evidence from US Patents.”

Kerr, William R., Ramana Nanda, and Matthew Rhodes-Kropf. 2014. "Entrepreneurship as Experimentation." *Journal of Economic Perspectives*, 28(3): 25–48.

Kueng, Lorenz, Mu-Jeung Yang, and Bryan Hong. 2014. "Sources of Firm Life-Cycle Dynamics: Differentiating Size vs. Age Effects."

Li, Guan-Cheng, Ronald Lai, Alexander D'Amour, David M. Doolin, Ye Sun, Vetle I. Torvik, Amy Z. Yu, and Lee Fleming. 2014. "Disambiguation and co-authorship networks of the U.S. patent inventor database (1975–2010)." *Research Policy*, 43(6): 941–955.

Li, Xiaoming, Guagquan Xu, Litao Jiao, Yinan Zhou, and Wei Yu. 2019. "Multi-layer network community detection model based on attributes and social interaction intensity." *Computers & Electrical Engineering*, 77: 300–313.

Magerman, Tom, Bart Van Looy, and Xiaoyan Song. 2006. "Data Production Methods for Harmonized Patent Statistics: Patentee Name Harmonization." KUL Working Paper MSI 0605.

Naidu, Suresh, Eric A. Posner, and Glen Weyl. 2018. "Antitrust Remedies for Labor Market Power." *Harvard Law Review*, , (2): 536–601.

Olmstead-Rumsey, Jane. 2019. "Market Concentration and the Productivity Slowdown." *MPRA Paper 93260.*

Park, Michael, Erin Leahey, and Russell J. Funk. 2023. "Papers and patents are becoming less disruptive over time." *Nature*, 613(7942): 138–144.

Peeters, Bert, Xiaoyan Song, Julie Caellert, Joris Grouwels, and Baart Van Looy. 2010. "Harmonizing harmonized patentee names: An exploratory assessment of top patentees." *Eurostat Working Paper.*

Piazza, Roberto, and Yu Zheng. 2019. "Innovate to Lead or Innovate to Prevail: When do Monopolistic Rents Induce Growth?" IMF Working Paper 19/294.

Poege, Felix, Dietmar Harhoff, Fabian Gaessler, and Stefano Baruffaldi. 2019. "Science quality and the value of inventions." *Science Advances*, 5(12).

Reinsdorf, Marshall B., David M. Byrne, and John G. Fernald. 2016. "Does the United States have a productivity slowdown or a measurement problem?" *Brookings Papers on Economic Activity*, 1: 109–182.

Rogerson, Richard, Robert Shimer, and Randall Wright. 2005. "Search-Theoretic Models of the Labor Market: A Survey." *Journal of Economic Literature*, 43(4): 959–988.

Sakakibara, Mariko, and Lee Branstetter. 2001. "Do Stronger Patents Induce More Innovation?" *RAND Journal of Economics*, 32(1): 77–100.

Scotchmer, Suzanne. 1991. "Standing on the Shoulders of Giants: Cumulative Research and the Patent Law." *Journal of Economic Perspectives*, 5(1): 29–41.

Silipo, Damiano B. 2005. "The Evolution of Cooperation in Patent Races:Theory and Experimental Evidence." *J Econ*, 85(1): 1–38.

Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the US Productivity Slowdown." *Journal of Economic Perspectives*, 31(2): 165–186.

Thompson, Neil C, and Jeffrey M Kuhn. 2020. "Does Winning a Patent Race lead to more follow-on Innovation?" *Journal of Legal Analysis*, 12: 183–220.

Toole, Andrew, Christina Jones, and Sarvothaman Madhavan. 2021. "PatentsView: An Open Data Platform to Advance Science and Technology Policy." USPTO Ec. WP 2021-1.

Zizzo, Daniel John. 2002. "Racing with uncertainty: a patent race experiment." *International Journal of Industrial Organization*, 20(6): 877–902.

Appendices

A Identifying Inventors in PATSTAT

Since PATSTAT does not contain IDs, only string names, I consolidate spelling mistakes and disambiguate entities with the same name before using the data. This appendix describes the procedure.

First, Magerman, Van Looy and Song (2006) have already constructed consolidated identifiers by correcting spelling mistakes, omitting titles and reading out abbreviations like "Ltd.". They have also constructed a sector variable, which assigns names in the database to categories like "company", "individual", "university" etc. After fusing such different spellings of the same name, they find an additional 30% of patents for the top 450 applicants, compared to the raw HAN identifiers provided by PATSTAT.

Second, Peeters et al. (2010) have manually checked the record of the top 450 applicants and searched for additional possible variants in the data. They can assign another 30% of patents to these applicants. However, since some of these applicants have over 100.000 patents in different countries, different spellings and mistakes play a much larger role than in the general population.

To disambiguate additional names both on the inventor and firm side, I clean names similarly to Magerman, Van Looy and Song (2006) and then sort all words alphabetically. This equates reversed spellings of names like "Erik van Houten" and "van Houten, Erik". This reduces the number of unique inventor identifiers by another 25%. I additionally clean firm names of addresses that are sporadically entered in the field "name", e.g. "Intel Corporation, Santa Clara, CA". This fuses around 3% of the remaining firm identifiers.

To gauge the quality of the resulting ID, I draw a list of prominent inventors from Wikipedia and link them to our data. Just as Peeters et al. (2010) for the firm side, I find that these highly active individuals are split over multiple IDs due to spelling mistakes, different name formats etc. However, the automated correction of Magerman, Van Looy and Song (2006) already does a decent job of aggregating them: After manual search, I e.g. link 38 PATSTAT person IDs to the most prolific inventor in the world (Dr. Shunpei Yamazaki). Magerman, Van Looy and Song (2006) already linked the most important 30, so I can only marginally improve upon their results. My 38 IDs participate in 5585 patent families across the world while the 30 IDs of Magerman, Van Looy and Song (2006) participate in 5581. The newly discovered name variants are clearly errors that only show up once. In addition, such spelling variants often show up within a patent family where the inventor is also cited on other patents. The patent family is the relevant unit of observation. Thus, even if undetected spelling variants exist, they are largely irrelevant to my productivity measures. I thus have confidence that the IDs provided by Magerman, Van Looy and Song (2006) capture the large majority of an inventor's patents.

However, this still leaves the problem that some names might belong to more than one inventor. Combining such inventors into one person would create the impression of a prolific inventor frequently moving between firms.

First, I collect the frequency with which words occur in the inventor names submitted on patents in each country. I then eliminate inventor names that do not contain two infrequent words: E.g., "Erik van Houten" contains two words common in Dutch names ("Erik" and "van") and only one uncommon word "Houten". Thus, I will not consider this inventor in the sample.

Second, PATSTAT contains the IPC classes associated with each inventor's patents. Inventors will typically not master a variety of technical fields and thus names with more diverse portfolios are more likely to stand for more than one inventor. Specifically, I exclude workers whose most common IPC 4-digit category accounts for 20% or less of their patents, whose top technology field accounts for 50% or less of their patents and whose top two technology fields account for 80% or less of their patents. I check these numbers against the statistics for inventors crosschecked with Wikipedia to guarantee that these criteria are not too strict.

Third, I exclude inventors from the sample who were active for more than 40 years, on the basis that these are likely overlapping inventors of the same name.

The observed time span, the diversity of IPC classes and technology communities and the number of distinct names are conceptually different criteria. Nonetheless, they are reasonably correlated (0.15-0.6), which suggests that the criteria identify suspect inventors reliably.

After these cleaning steps, I obtain inventor IDs that connect all patents strings that could conceivably belong to the same inventor, but could also belong to different persons. I feed the resulting IDs into the disambiguation algorithm of Li et al. (2014). This algorithm constructs a similarity score for different records belonging to the same last name and connects those that are most likely from the same person, based on coauthors, coapplicants, geographic proximity and the abstract of each patent.

B Model Extension: Vacancies and Unemployment

New inventors search for firms' open vacancies. In contrast to the standard search and matching labor market model, I assume that new inventors enter the labor market, immediately find a match among the available vacancies and that unmatched inventors have to leave the economy because they lose their connection to recent developments. The research avenues that are represented by vacancies also become superseded by new approaches if they do not match. This reduces the complexity of the labor market, because the mass of unemployed inventors does not matter for the equilibrium anymore, since they cannot contribute to the economy. This simplifying assumption eliminates two state variables from each field's inventor labor market: the number of unmatched vacancies and the number of unmatched inventors. This assumption leads to the same steady state outcome, but the path towards that steady state is much more tractable. Figure (8) describes the path towards labor market equilibrium after a disruptive innovation for both specifications.

Figure 8: Simulated Example Labor Market with Vacancies

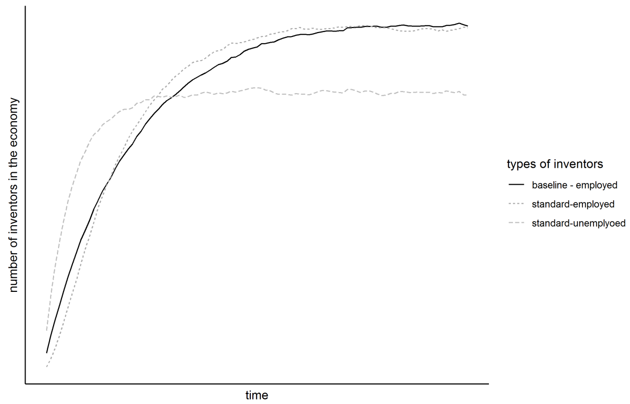

Notes: The graph shows the evolution of the number of incremental inventors in a technology cluster after its foundation. Over time, more and more inventors enter the cluster, until the steady state level is reached. The baseline specification of the model is presented in black. The grey lines depict the stock of employed and unemployed inventors in a more standard model for comparison. Such a model has slightly less employed inventors early on, because inventors enter into unemployment and leave it over time. However, not only do both models give the same kind of path qualitatively, the two paths are also quantitatively close. Assuming that inventors cannot be unemployed increases tractability without greatly changing even the quantitative results. *Sources*: Own simulations.

How many vacancies firms will create in this setting depends on the value of obtaining an additional inventor. This value is determined by the number of patents the new inventor will produce and by how much the firm has to pay the inventor.