

DISCUSSION PAPER SERIES

DP17911

INNOVATION NETWORKS AND BUSINESS-STEALING

Philippe Aghion, Matthew O. Jackson, Antoine
Mayerowitz and Abhijit Tagade

**MACROECONOMICS AND GROWTH
AND ORGANIZATIONAL ECONOMICS**



INNOVATION NETWORKS AND BUSINESS-STEALING

Philippe Aghion, Matthew O. Jackson, Antoine Mayerowitz and Abhijit Tagade

Discussion Paper DP17911
Published 16 February 2023
Submitted 20 December 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Macroeconomics and Growth
- Organizational Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Philippe Aghion, Matthew O. Jackson, Antoine Mayerowitz and Abhijit Tagade

INNOVATION NETWORKS AND BUSINESS-STEALING

Abstract

We use the universe of US Patent and Trademark Office (USPTO) data on patents and inventors from 1976 to 2017 to look at how inventors' potential concern for business-stealing affects coauthorship on patents. First, we find an inverted-U shape in the fraction of coauthors that an inventor has per year who are new as a function number of other inventors also working in an inventor's field. Second, we find that after a breakthrough invention, an inventor brings in persistently fewer than usual new coauthors. Third, a higher potential concern for business stealing—as measured either by the number of others working or the average price markups by firms in the area—leads to a higher drop in the fraction of new co-authors per patent after a breakthrough. We show how these patterns can be explained via a simple model in which inventors trade off gains from collaboration against threats of business stealing.

JEL Classification: D85, L14

Keywords: Innovation, Business stealing, Collaboration

Philippe Aghion - philippe.aghion@college-de-france.fr
Collège de France, London School of Economics, INSEAD and CEPR

Matthew O. Jackson - jacksonm@stanford.edu
Stanford University

Antoine Mayerowitz - a.mayerowitz@gmail.com
Paris School of Economics (PSE)

Abhijit Tagade - a.tagade@lse.ac.uk
_ *CEPR Admin team*

Innovation Networks and Business-Stealing*

Philippe Aghion Matthew O. Jackson

Antoine Mayerowitz Abhijit Tagade[†]

Abstract

We use the universe of US Patent and Trademark Office (USPTO) data on patents and inventors from 1976 to 2017 to look at how inventors potential concern for business-stealing affects coauthorship on patents. First, we find an inverted-U shape in the fraction of coauthors that an inventor has per year who are new as a function number of other inventors also working in an inventor's field. Second, we find that after a breakthrough invention, an inventor brings in persistently fewer than usual new coauthors. Third, a higher potential concern for business stealing—as measured either by the number of others working or the average price markups by firms in the area—leads to a higher drop in the fraction of new co-authors per patent after a breakthrough. We show how these patterns can be explained via a simple model in which inventors trade off gains from collaboration against threats of business stealing.

Keywords: Innovation, Business Stealing, Collaboration, Knowledge Networks

JEL Classification Codes: D85, L14

*We thank Xavier Jaravel for many helpful discussions and John Van Reenen for helpful comments. We gratefully acknowledge financial support under NSF grant SES-2018554.

[†]Aghion is at the College de France, London School of Economics, and INSEAD; Jackson is at the Department of Economics, Stanford University, and is an external faculty member at the Santa Fe Institute; Mayerowitz is at Paris School of Economics and the College de France; Tagade is at the Department of Economics, London School of Economics. Emails: philippe.aghion@college-de-france.fr, jacksonm@stanford.edu, antoine.mayerowitz@college-de-france.fr, and a.tagade@lse.ac.uk.

1 Introduction

Research involves putting ideas together and building upon knowledge held by different people and sources. There is substantial evidence that teams can be more effective than individuals and that connections between groups help spark new ideas (e.g., Walker et al. 1997; Burt 2004; Hong et al. 2004; Page 2007; Ozman 2009, etc.), with larger teams leading to more breakthroughs and creativity. Although such positive externalities and synergies offer incentives for researchers and inventors to collaborate, inventors also face forces that push in the opposite direction. In particular, they may face competition from their collaborators as well as broader leakage of information about their inventions. In this paper, we examine the implications of these countervailing forces both empirically and theoretically.

For our empirical analysis, we use the universe of the US Patent and Trademark Office (USPTO) data on granted patents and inventors from 1976 to 2017 to look at how inventors' potential concern for business-stealing affects coauthorship on patents. First, we find an inverted-U relationship in the growth in coauthorship as a function of the number of other inventors working in an area: the rate per year at which inventors add new collaborators is highest at intermediate levels of competition in terms of the number of other inventors patenting in the same subject area. Second, we use an event study method to examine how networks are influenced by having a major breakthrough, which we measure as a patent in the top 2 percentiles of citations over following years. We find that after a breakthrough, the fraction of coauthors who are new on the inventor's patents drops significantly below what it was before the breakthrough and does not recover to its pre-breakthrough value. Third, we examine how the extent to which an inventor stops working with new coauthors depends on two measures of potential losses from business stealing: the average price markups by firms (aggregated at the level of NBER technology sectors), as well as the number of other inventors in the area. In both cases, there is a significantly higher drop in the number of new co-authors after the breakthrough when there are greater potential losses from business stealing as measured by either the number of competitors or price markups.

We also provide a model that offers an explanation as to why these effects should be expected; and the simplicity of the model shows how robust these effects can be. The two forces in the model are countervailing: having greater numbers of collaborators on a team increases the probability of having a successful invention, but having more new

collaborators can lead to lower potential profits for the inventor from the breakthrough, both by spreading profits more thinly and by increasing the probability that information leaks to others in the field. Moreover, the potential loss in profits due to information leakage and spreading the profits among more total team members are greater with greater numbers of potential competitors and/or higher markups. The results from the model are consistent with the patterns that we discover in the data.

The implications of our findings are important for economic efficiency and growth. Fears of business stealing and leakage of ideas lead to lower investment in collaborative research, and our analysis shows that these effects can be substantial and hold across patent areas, especially those with many inventors or high price markups.

Our paper relates to several strands of literature. First, we add new insights into the literature on networks and how communication depends on worries of sharing information (e.g., Immorlica et al. 2014). Indeed, we are not the first to recognize that business stealing can lead to inefficiencies. Stein (2008) shows how information exchange can be hampered when players are in competition with each other, and that can be exacerbated when information can diffuse. Dasaratha (2019) examines a game of network formation among competitors who can share ideas to innovate, but those might be then shared further in the network. He finds that the network ends up being exactly at a critical threshold between sparse and dense networks, and shows that welfare can be improved by having public innovators who freely share information. Our contribution is to provide evidence of the importance of business-stealing in the dynamics of research networks, and to show and model how the fear of business-stealing interacts with competition in the inventor's field.

Also related to our analysis is the literature on research teams and research output. For example, Akcigit et al. (2018) show how interacting with other researchers, particularly with top researchers, positively affects innovation. Azoulay et al. (2010) and Jaravel et al. (2018) show how the death of a top innovator negatively affects the subsequent productivity of their coauthors. We contribute to these literatures by pointing at the importance of business-stealing and competition in the dynamics of network building.

Our paper also relates to the literature on competition and innovation. In particular, Aghion et al. (2005) describe an inverted-U relationship between innovation and competition. Our finding instead is about the formation of teams and the inverted-U is between an inventor's fraction of new collaborators per year and the potential losses from business stealing. Thus, our inverted-U is between different things and occurs for

different reasons. Moreover, our event studies have no parallel in the previous literature. The patterns that we uncover on breakthrough patents leading to fewer than usual new collaborators, and more extremely with more potential losses from business stealing, are novel and provide interesting facts and questions for further research.

The remaining part of the paper is organized as follows. Section 2 presents our data and measures, Section 3 presents our empirical results, Section 4 presents our model, and Section 5 concludes.

2 Data, measurement and empirical strategy

2.1 Data description

We use the universe of USPTO data on inventors and utility patents that are filed from 1976 to 2017 and eventually granted. Throughout the paper, the year of patent refers to its year of application, and a patent’s citations are measured within its first five years of application.

We identify inventors and assignees using the disambiguation algorithm of Monath et al. (2015) available on the *Patentsview* database (our data includes assignees and inventors based both inside and outside the US).

We use the coauthorship of inventors to build collaboration networks, and track them over time. Similarly, patents are connected to each other through the citations that appear in the patents.

A statistical overview of these networks is presented in Table 1. Inventors in more recent years have more than three times as many coauthors as they did in the 1970s and 80s. Similarly, patents have become more interconnected over time, on average.

Some coauthors on patents, such as intellectual property lawyers, management executives, etc., may not contribute to the creative process of an invention. To filter out such individuals, we restrict our attention to those inventors who within our data: (i) publish at least one patent alone, (ii) publish at least three patents, (iii) have a research career that lasted more than three years, and (iv) coauthor with at least three other people in our data.

We present the summary statistics for the selected and the full population of inventors in Table 2. Overall, our selected sample covers about 25 percent of the people in our data,

Table 1: Descriptive statistics of coauthorships and citations networks of the full data

<i>Coauthorship Network</i>	1976 - 1985	1986 - 1995	1996 - 2005	2006 - 2015
Number of inventors	522,279	845,907	1,396,121	1,481,191
Average number of coauthors	3.63	5.83	9.06	11.60
Giant component size	138,116	382,120	858,732	992,365
Overall clustering	0.76	0.77	0.76	0.77
<i>Citation Network</i>	1976 - 1985	1986 - 1995	1996 - 2005	2006 - 2015
Number of patents	592,789	1,012,254	1,934,934	2,012,732
Average citations	1.95	3.22	5.66	4.14
Giant component size	548,503	988,623	1,914,253	1,873,129
Overall clustering	0.10	0.10	0.10	0.14

Notes: Descriptive statistics of the coauthorship and citation networks over sequential periods of time using Patentsview data. *Giant component size* refers to the number of nodes in the largest connected component. *Overall clustering* examines situations in which one node is connected to two others and reports the fraction of those for which the two others are connected to each other.

and those in our sample tend to have a broader network of coauthors, longer experience, and more published patents.

2.2 Measuring changes in inventors' networks

To analyze the evolution of an inventor's network, we build a panel of inventors and track their collaborations over time. We consider two measures of how inventors expand their networks.

Our first measure of the expansion of an inventor's network is based on the fraction of collaborators who are new in each year. Let $C_{i,t}$ denote the number of collaborators of an inventor i during year t , and $S_{i,t}$ be their stock of (unique) collaborators up through year t . The fraction of new collaborators is:

$$f_{i,t} = (S_{i,t} - S_{i,t-1}) / C_{i,t}. \quad (1)$$

This measure is normalized per number of collaborators in a given year to account for the fact that the stock is growing over time, and so we divide by C and not S . Since collaborators of an inventor are observed only in years during which they apply for a patent, the value of f is imputed to the last observed value during the years in which no

Table 2: Cross sectional summary statistics of the full and selected inventors' data

Year of application:	1985			1995			2005		
Variables	All	Selected	Treated	All	Selected	Treated	All	Selected	Treated
Degree									
Mean	14	15	22	23	27	42	38	48	77
Median	6	8	10	10	12	20	16	24	42
Growth rate of degree	0.11	0.12	0.14	0.15	0.14	0.16	0.13	0.12	0.13
Fraction of new coauthors	0.79	0.74	0.72	0.73	0.69	0.64	0.69	0.65	0.59
Experience	5.98	6.77	6.92	8.7	10.2	10.9	10	13	14
Patents per inventor	0.47	0.55	0.74	0.64	0.77	1.10	0.69	0.88	1.26
Team size	2.59	2.47	2.64	3.19	2.93	3.20	3.71	3.08	3.33
Patenting inventors	32%	33%	39%	38%	39%	48%	38%	44%	53%
CPC Section									
Chemistry and metallurgy	21%	18%	19%	19%	14%	15%	16%	11%	12%
Electricity	16%	19%	19%	19%	22%	23%	24%	26%	27%
Fixed constructions	3.1%	2.8%	2.2%	2.4%	2.3%	1.9%	2.0%	2.0%	1.9%
Human necessities	11%	9.5%	8.5%	12%	10%	9.7%	12%	9.9%	9.5%
Mechanical engineering	9.9%	9.9%	8.5%	8.1%	8.4%	7.7%	6.6%	7.5%	7.4%
Performing ops.; Transporting	18%	18%	17%	16%	16%	15%	13%	14%	13%
Physics	19%	21%	24%	22%	25%	26%	27%	29%	29%
Textiles and paper	2.1%	2.2%	1.5%	1.5%	1.5%	1.1%	0.8%	0.8%	0.7%
Country									
Other	50%	48%	47%	51%	51%	46%	51%	54%	46%
US	50%	52%	53%	49%	49%	54%	49%	46%	54%
Assignee type									
Firm	89%	94%	96%	92%	95%	97%	95%	97%	97%
Other	11%	6.2%	4.1%	7.8%	4.8%	3.2%	4.7%	3.4%	2.7%
Number of inventors	167,989	41,863	7,382	309,519	74,980	13,574	446,534	93,514	17,244

Notes: Summary statistics of the full and selected panel data of inventors for three given years. Our selected sample of inventors are those who have at least one single-authored patent, and a career spanning at least 3 years, with 3 granted patents, and having at least 3 coauthors within our data. “Treated” inventors (in our event study) are the selected inventors that publish one breakthrough patent over their career. A breakthrough is a patent that is among the top 2 percent most cited patents during the five years following its application, among the pool of all patents filed within two years of the patent in its CPC subclass. We discuss this further in section 2.4. *Degree* is defined as an inventor’s cumulative number of unique coauthors up through a given year. *Growth rate of degree* is the growth rate of inventor’s *degree*. *Fraction of new coauthors* is the fraction of an inventor’s coauthors of patents in a given year who are new. *Experience* is the number of years since the filing date of the inventor’s first granted patent. *Patents per inventor* are the average number of patents filed per inventor in the given year (that were eventually granted). *Team size* is the number of coauthors on the last filed patent. *Patenting inventors* are the share of inventors applying for a patent in that year. *CPC Section* shows the distribution of inventors in each broad technology class as defined by the USPTO.

patent was filed.

Our second measure of the expansion of an inventor’s network is the growth rate of the stock of coauthors as measured in the change of the log:

$$\Delta \log S_{i,t} = \log(S_{i,t}/S_{i,t-1}). \quad (2)$$

Since $S_{i,t}$ does not change until an inventor collaborates with new coauthors on patents, this measure is 0 until the inventor collaborates with a new coauthor.

Given that our main focus is on inventors’ willingness to expand their network subject to the threat of business-stealing, we consider the first measure to be preferable for our empirical exercises, since it is robust to the productivity of inventors. By contrast, although the growth rate of new coauthors allows us to estimate the impact on aggregate networks of inventors, it depends on research productivity. Nonetheless, it is a useful robustness check and alternative view of the evolution of the network.

2.3 Measuring the potential concern for business-stealing

We consider several measures of the potential concern for business-stealing, each capturing a different aspect of potential concerns.

First, areas with greater numbers of inventors reflect a greater potential competition for profits from innovation than those areas with fewer inventors, and can also reflect a greater potential possibility of leakage of information. Thus as our most basic measure of potential concerns of business stealing, we consider the number of active inventors in a USPTO CPC subclass.¹ An inventor’s CPC is defined by the technology they most frequently publish patents in. The number of active inventors in a given year and CPC subclass are counted as those who have patented at least once in that subclass by the end of the year, and will patent again in the future in that subclass. The distribution of number of inventors in CPC subclasses is presented in Figure 6 in the Appendix. Since inventive activity in technologies can change in the long run, we re-classify the CPC subclasses into high or low concern for business stealing every year.

Second, we also consider the rate of entry of new inventors, classifying those tech-

¹A subclass of a patent in our data is a four character USPTO CPC classification. For example, *H04N* is a subclass that represents the technology of *Pictorial Communication*.

nologies with more inventors entering per year as being ones with a higher level of competition. Entry rate is estimated as the ratio of the number of inventors who are patenting for the first time, to the number of patenting inventors in a given year in each technology category. Because the number of new inventors patenting in various CPC subclasses each year may be too small to provide sufficient precision for our analysis, we use the broad NBER technology categories as our unit of aggregation for the entry rate of inventors.² We treat inventors patenting in the NBER categories with an above median entry rate as those facing a higher than average concern of business-stealing.³

Third, we examine average size of price markups in the corresponding sector: higher markups can lead to higher potential lost profits from business-stealing. We estimate markups for each NBER technology sector using a three step process: first, we measure markups at the firm level using the method of De Loecker et al. (2020) using the Compustat database. Second, we match firm level markups with patent assignees using the data of Dorn et al. (2020). Third, weighting by patent citations, we estimate the markups for each NBER technology sector using the matched firm level markups, and update its value every three years.⁴

Fourth, we construct a localized measure of competition for each inventor based on how much competition each of their patents faces. For this purpose, we build a network of *similar* patents and estimate the density of this network. This captures how intensely researchers are working on similar problems. In particular, say that two patents are *similar* to one another if the patents they they cite overlap. We estimate the Jaccard index (size of the intersection over union) of the backward citations between all pairs of patents, and consider that two patents (vertices) are connected by an undirected and unweighted edge if their similarity is greater than 5%.⁵ Since two patents published many years apart are less likely to be in competition than if they are published within a few years of each other, we retain only those edges between similar patents that are published within 3 calendar years of each other. We also use this narrow time band to have a measure that

²See Marco et al. (2015) for details on patent classification by NBER technology categories.

³Replicating our empirical exercises by aggregating entry rate of inventors at the level of CPC class does not change the direction of our results.

⁴Repeating our estimations with alternative weightings—by patent applications and firm revenue—we find no significant changes in our results.

⁵For a pair of patents A and B, their respective set of backward citations (who they cite) C_A, C_B and the set of technologies—measured by the CPCs—of their backward citations $T_i = T(C_i) \ i \in \{A, B\}$, the similarity index is defined as $S(A, B) = S(B, A) = |T_A \cap T_B| / |T_A \cup T_B|$.

is less likely to be contaminated by knowledge spillover.⁶ Within this network of patents, we identify patent communities using the Louvain algorithm developed by Blondel et al. (2008) and estimate a community’s density as its total number of connections normalized to the total possible connections within the community. More formally, we define the density measure of community C of similar parents as:

$$dens(C) = \left(\sum_{\{p,q\} \subset C, p \neq q} \mathbb{1}_{p,q} \right) / \binom{N(C)}{2}$$

where $N(C)$ is the number of patents in community C and where $\mathbb{1}_{p,q}$ is equal to 1 if the two patents p and q are connected, and is equal to zero otherwise. Since density can be correlated with the size of a community, we residualize the density measure $dens(C)$ by the size of C ; i.e., by the number of patents in C . Note that some patents are isolated: without other patents in their community. In that case, we define their density as 0, meaning that such patents do not face any competition given no patent is similar to them.⁷

2.4 Breakthrough inventions and event studies

We examine how a breakthrough invention impacts a inventor’s collaboration, and how this effect depends upon the measures for potential concern of business stealing listed above.

We define a *breakthrough* as a patent that is among the top 2 percent most cited patents within five years following its application, among the pool of all patents applied for within two years in its CPC subclass. The last granted patent in our dataset is applied for in 2015, and its citations are measured up to (including) 2017. To minimize the chance that we misidentify a patent as a breakthrough, we only consider those CPC subclasses with at least 100 patents in their comparison set.

The “event” we consider is the year of application of a patent that is a breakthrough, and our treated inventors are the ones who author them. We time the event by the year

⁶Distinguishing business stealing from knowledge spillovers is challenging (e.g., see the discussion in Bloom et al. 2013), and our results should be interpreted with corresponding caution.

⁷Some inventors patent in multiple research areas and may face different levels of competition in each of them. We define the degree of competition faced by an inventor to be the average of the level of competition faced by each of their patents in their corresponding communities.

of patent’s application. We compare the changes in the inventor’s network normalized to the stream of previous patents that were not breakthroughs, as a proxy for what would have happened if no breakthrough had occurred.

In our baseline event study exercise, we look at how a breakthrough patent affects the growth rate of an inventor i ’s network at period t as captured by their fraction of coauthors that are new connections (i.e., $f_{i,t}$), as well as the growth rate of their stock of coauthors (i.e., $\log(S_{i,t}/S_{i,t-1})$). We then decompose the baseline by looking at how a breakthrough invention’s impact on the dynamics of coauthorship depends upon the threat of business-stealing.

We restrict our analysis to inventors who have only one breakthrough over their research career to capture the effect of business stealing on networks in our event study. A statistical overview of inventors with a breakthrough is presented in the *treated* columns in Table 2.

3 Empirical analysis of the dynamics of collaborations

3.1 Cross sectional analysis

Looking at all inventors, we first examine how our main measure of collaboration of inventors—the fraction of coauthors who are new connections in a given year—depends on our basic measure of concerns of business-stealing—the density of active inventors in a patent’s subclass.

Figure 1 and Table 3 show an inverted-U relationship between the density of the technological sector and the willingness of inventors to add new collaborators to their team. We find positive and significant regression coefficients on concerns of business-stealing (density), and negative and significant regression coefficients on the square term. Controlling for variation both by year, and by technology classes within a year, the direction of their relationship remains stable. Controlling for technology subclass in addition to year, the fraction of new coauthorships is (weakly) negatively correlated with the density of inventors.

Intuitively, starting from zero density, higher density increases the scope for coauthorship as there are more researchers working in an area, hence the upward part of the inverted-U; however too much density induces the inventor to reduce coauthorship so

Table 3: New collaborations as a function of the density of inventors

	Fraction of coauthors who are new			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
(Intercept)	0.612*** (0.017)			
Log number of inventors	0.049*** (0.006)	0.054*** (0.005)	0.035*** (0.005)	-0.015* (0.008)
Log number of inventors squared	-0.005*** (0.000)	-0.005*** (0.000)	-0.003*** (0.000)	-0.002*** (0.001)
<i>Fixed-effects</i>				
Year		✓		✓
Year × CPC class			✓	
CPC subclass				✓
<i>Fit statistics</i>				
Observations	17,842	17,842	17,842	17,842
R ²	0.02042	0.40238	0.62621	0.63106
Within R ²		0.01104	0.00507	0.02623

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: The table shows the estimates of the regression of the average fraction of coauthors who are new connections in a given year for patenting inventors within a CPC subclass, on the log number of inventors in the CPC subclass. Standard errors are in parenthesis. *CPC class* is one unit of patent classification above CPC subclass.




Figure 1: Correlation Between the Density of Inventors and New Collaborations

Notes: The figure plots the binned scatter relationship between the average fraction of coauthors in a given year who are new connections for its patenting inventors (along the vertical axis) and the log of the number of active inventors in a CPC subclass (along the horizontal axis) using 100 bins. The unit of observation is a CPC subclass. The solid line is the least squares fit with a second order polynomial. Both variables are residualized by year and centered around the sample average.

as to limit the risk of business/knowledge-stealing. Absent any risk of business-stealing, one would expect the relationship between new collaborations and people working in the area to be increasing, or at least nondecreasing.

3.2 An event study of breakthrough patents

Next, we measure the pattern of inventors' coauthorships during and after the filing of an application for a breakthrough patent. This gives us a different perspective, and allows us to see if inventors act differently once they make a big discovery.

The idea is that inventors should be more willing to collaborate when they are looking for breakthrough ideas than once they have one in hand. Once they have a breakthrough there is less value from adding new collaborators and more potential risk in losing profits and having information leak.

Our identification strategy relies on a difference-in-differences design with staggered adoption. The empirical model for the outcome of interest for an inventor i on year t is:

$$Y_{i,t} = \alpha_i + \lambda_t + b_{i,t} \times Treated_{i,t} + \varepsilon_{i,t} \quad (3)$$

where α_i and λ_t capture inventor and year fixed effects, respectively, on the dependent variable of inventor network measured by $Y_{i,t}$, and $Treated_{i,t} = 1[t \geq E_i]$ is a dummy variable indicating that the inventor is treated, where E_i is the year where inventor i makes an application for her breakthrough patent. Finally, $b_{i,t}$ is the individual treatment effect, that is, the effect of having a breakthrough patent. Model (3) works from a parallel trends assumption: had there been no breakthrough patent, the expected outcome of an inventor would be $\alpha_i + \lambda_t$.

The canonical method to estimate the treatment effect in this setting is an Ordinary Least Squares (OLS) with two-way fixed effects (TWFE) and some leads to test for pre-trend along with lags to recover dynamic effect. However recent literature has shown that this estimator is not reliable in the presence of treatment effect heterogeneity.⁸ We thus use the *imputation estimator* proposed by Borusyak et al. (2021) for its ease of interpretation, efficiency, and performance with large sample.

The estimation is a three step procedure. First, inventor and year fixed-effects α_i, λ_t in (3) are estimated on not yet treated observations, that is observations where $Treated_{i,t} = 0$. Second, we compute the treatment effect $\hat{b}_{i,t} = Y_{i,t} - \hat{\alpha}_i - \hat{\lambda}_t$ of each treated observations using the estimated fixed effect. Finally, averaging $\hat{b}_{i,t}$ gives a consistent estimate of the average treatment effect on the treated β . The procedure allows estimation of the dynamic effect $\hat{b}_h$ for each horizon h relative to the treatment year by averaging individual treatment effect $\hat{b}_{i,t}$ at each horizon. That is

$$\hat{b}_h = \frac{1}{|I_h|} \sum_{i \in I_h} \hat{b}_{i,t} \quad (4)$$

with I_h the set of inventors observed h periods after treatment. We focus on $h \in \{0, \dots, 10\}$.

To provide support for our empirical analysis, we also test for parallel trends following Borusyak et al. (2021). We run a TWFE regression on not yet treated observations:

$$Y_{i,t|t < E_i} = \alpha_i + \lambda_t + \sum_{p=-P}^{-1} \gamma_p Treated[t = E_i + p] + \varepsilon_{i,t} \quad (5)$$

where E_i is the year of treatment and $Treated[t = E_i + p]$ is an indicator variable of

⁸See De Chaisemartin et al. (2020), Borusyak et al. (2021) and Goodman-Bacon (2021).

being treated p years later. In our event study plots, we combine both treatment effect coefficients $\{\hat{b}_h\}$ and pre-trend coefficients $\{\hat{\gamma}_p\}$. In all subsequent analysis, we report the standard errors clustered at the inventor level.⁹

In our baseline regression, we adopt a richer model that allows time fixed effects to vary by technological field at the CPC subclass level and by the inventor’s country of origin. We also control for a life-cycle effect by adding the experience of the inventor as measured by the number of years since the application of their first patent observed in our dataset.¹⁰

$$Y_{i,t} = \alpha_i + (Year \times CPC\ subclass \times Country)_{i,t} + Experience_{i,t} + b_{i,t} \times Treated_{i,t} + \varepsilon_{i,t} \quad (6)$$

We conduct event studies using this approach for each of our two network measurements, namely fraction of new coauthors as well as the growth rate of the stock of coauthors.

3.3 Results

Figure 2 (see also 7) and Table 4 show the results of our baseline event study exercise. The bars in the figure are 95% confidence intervals.

We see a large and significant drop in the fraction of new coauthors per year in the year of the event and subsequent.¹¹ Treated inventors work with fewer than usual new coauthors on research teams during and subsequent to a breakthrough, and do not return to their previous level.

⁹While it would be even more desirable, clustering at team level is not tractable because teams change over time.

¹⁰Since the fully tracked USPTO PatentsView data series begins in 1976, we are unable to retrieve the true years of experience for inventors active prior to 1976. Therefore, inventors with 2 years of observed experience as of 1978, say, may have more years of experience. This may bias our event study estimates, particularly for inventors observed early in their career. However, we do not find reason to believe that they are differentially biased by concerns for business-stealing. Moreover, limiting our data to inventors who are first observed in years beginning 1986 onward, which keeps only those few inventors active prior to 1976 that did not patent in over 10 years, does not change our results.

¹¹The drop in the year of the breakthrough is smaller than subsequent years. There can be other patents in that year that come before the breakthrough, and it might also take some time to realize that one has a breakthrough. In 7 we actually see a jump up at the time of a breakthrough. That reflects the fact that the year of a breakthrough has a selection bias to have more patents than usual and hence more new coauthors than in a typical year. That is controlled for in fraction of new coauthors calculation, as that controls per patent, and so this effect is eliminated in Figure 2.




Figure 2: Event study of an application for a breakthrough patent on fraction of coauthors on patents who are new connections

Notes: The figure plots the treatment effect of a breakthrough patent on inventor collaborations using the Borusyak et al. (2021) imputation estimator. The event is the application for a patent by an inventor that is among the 2 percent most cited patents filed in its CPC subclass within a range of 2 years; the event is timed by the year of its application. The outcome is the fraction of new coauthors among all coauthors in a given year of an inventor. Pre-trends are estimated using OLS on not-yet-treated observations. For years with no patent, the value is imputed using the last observed value. Bars indicate the 95% confidence interval. The estimation controls for inventor, experience, and for year×CPC subclass×country specific shocks.

Table 4: Baseline effect of a breakthrough on subsequent inventor coauthorship

	Fraction of coauthors who are new				Growth rate in stock of coauthors			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated	-0.194*** (0.003)	-0.081*** (0.005)	-0.072*** (0.005)	-0.152*** (0.006)	-0.123*** (0.003)	-0.058*** (0.003)	-0.052*** (0.004)	-0.071*** (0.004)
Inventor	✓	✓	✓	✓	✓	✓	✓	✓
Year		✓				✓		
Year×CPC×Country			✓	✓			✓	✓
Experience				✓				✓
Wald stat.	137.139	10.470	4.305	1.055	127.816	4.346	1.591	2.030
Observations	349130	343816	326066	326066	349130	343816	326066	326066
Units	18899	18899	18899	18899	18899	18899	18899	18899

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table summarizes the change in the inventor networks, as measured by our two metrics, subsequent to the application of a breakthrough patent, relative to had they not. We note that the coefficients are negative regardless of the controls, and the effect size remains similar controlling for year of application, the CPC subclass of the applied patent, and country of origin of the inventor. The coefficients for columns (1) to (4) are interpreted as the change in growth rate of stock of coauthors of an inventor, and for columns (5) to (8) as the percentage point change in team members that are new coauthors, subsequent to application of a breakthrough patent relative to inventors who have not yet applied for or published a breakthrough patent. Standard errors are clustered at the inventor level.

The baseline effect sizes summarized in Table 4 show that the measured change in coauthorship with new inventors is negative, and statistically significant, regardless of the granularity of controls.¹²

One alternative explanation for why successful inventors end up reducing coauthorship after a breakthrough, is that they enter in a different phase of production, where instead of inventing (exploring) they turn to developing the product (exploitation).

If that is the case, then the extent to which successful inventors reduce new collaborations would no reason to vary as a function of potential business-stealing. To explore this, we interact the breakthrough event with the degree of concern for business-stealing faced by the inventor.

Following our previous notation, we define $b_{h,q}$ as the treatment effect of having a breakthrough during a time horizon h , and belonging to a level of concern for business-stealing $q \in \{\underline{q}, \bar{q}\}$, which we denote as being above median in $\bar{q}$. This is a natural extension of equation (4), and the treatment effect is given as the average value of $b_{i,t}$ for inventors within each group. That is,

$$\hat{b}_{h,q} = \frac{1}{|I_{h,q}|} \sum_{i \in I_{h,q}} \hat{b}_{i,t} \quad (4')$$

where $I_{h,q}$ is the set of inventors belonging to group q and observed h periods after treatment.

Table 5 presents the results of the event studies decomposed by concerns for business-stealing using our four different measures of potential business stealing concerns (as defined in Section 2.3). The event studies for price markups and the entry rate of new inventors are presented in Figures 3a and 3b.

First, note that the baseline event study estimates hold for both high and low levels of concerns: breakthroughs induce a decline both in the fraction of coauthors who are new and in growth rate of the stock of coauthors. Second, this decline in is significantly larger among inventors that face a higher threat of business-stealing by every one of the measures. Moreover, the difference in the decline between above and below median

¹²Note that team members on a given breakthrough patent may differ in their measured effect size due to differences in baseline rates of new coauthorship, previous coauthorship with team members, and stage of the inventor's research career. Our results indicate that controlling for experience, as measured by years elapsed since the application of first patent, magnifies the effect size.

Table 5: Drops in new coauthors as a function of measures of concerns for business-stealing

	Fraction of coauthors who are new				Growth rate in stock of coauthors			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Density of Inventors: < median	-0.135*** (0.009)				-0.051*** (0.006)			
Density of Inventors: > median	-0.162*** (0.007)				-0.083*** (0.005)			
Sector Markup: < median		-0.123*** (0.009)				-0.045*** (0.007)		
Sector Markup: > median		-0.165*** (0.007)				-0.082*** (0.004)		
Entry Rate: < median			-0.104*** (0.007)				-0.056*** (0.005)	
Entry Rate: > median			-0.203*** (0.007)				-0.083*** (0.005)	
Similarity Density: < median				-0.149*** (0.007)				-0.038*** (0.005)
Similarity Density: > median				-0.163*** (0.008)				-0.106*** (0.005)
Inventor	✓	✓	✓	✓	✓	✓	✓	✓
Year×CPC×Country	✓	✓	✓	✓	✓	✓	✓	✓
Experience	✓	✓	✓	✓	✓	✓	✓	✓
Wald stat.	0.811	1.020	1.335	1.043	1.246	1.705	1.562	5.105
Z score	2.337	3.713	10.000	1.324	4.005	4.479	4.209	9.176
Observations	326064	322549	321027	326066	326064	322549	321027	326066
Units	18899	18872	18803	18899	18899	18872	18803	18899

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table summarizes the event study estimates for the effect of a breakthrough on inventor networks decomposed by the various measures of concerns for business-stealing discussed in section 2.3. We estimate the coefficient for two levels of each measure: above median reflects a higher concern for business-stealing over below median. Coefficients in columns (1) and (5) correspond to the density of inventors in their main CPC subclass; columns (2) and (6) to the entry rate of inventors in the different NBER categories; columns (3) and (7) to the measure of sector level markup faced by inventors; and, columns (4) and (8) to our measure pertaining the density of patent similarity in inventors' areas of research. The wald statistic is testing for nullity of 10 pre-trend coefficients using an OLS regression on not-yet-treated observations. Z score is computed assuming the absence of covariance between the two parameter estimates. Standard errors are clustered at the inventor level.

measures is statistically significant for each specification. Depending on how we measure concern for business-stealing, this additional drop amounts to between 1.4 to about 10 percentage points in the fraction of coauthors who are new connections, and between 2.7 to about 6.8 percentage points in the growth rate of in the stock of unique coauthors.

Our results are consistent with business-stealing playing an important role in inventors' decisions to (not) expand their coauthorship networks. The difference between the drop across these different measures could conceivably be correlated with some other unmeasured variables that differ from business stealing and lead team structures to diverge, but the fact that this holds for all four measures makes such alternatives less plausible.

4 A business-stealing explanation for the empirical results

Next, we propose a model that illuminates the empirical findings. It provides a direct explanation for: the inverted-U relationship; the fact that the growth rate in the number of collaborators drops after a breakthrough innovation; and the finding that the negative impact of the breakthrough invention on the subsequent growth in coauthorship is more pronounced the higher inventors concern for business-stealing. It is one possible model with these features, concentrating on the obvious tradeoff between an increased probability of breakthroughs from adding more collaborators versus decreased profits for the inventor conditional on a breakthrough.

4.1 A Base Static Model


We begin with a simple static model.

An inventor has access to up to n collaborators.


The inventor chooses a number d of collaborators.

The project has a probability of $P(d)$ of being successful, where d is the number of collaborators (beyond the inventor) who work on it. $P(0) > 0$ and $P(\cdot)$ is increasing and strictly concave in d .

If successful, the project yields profits $\pi_b(d)$. The subscript b refers to the level of the threat of business stealing. We take $\pi_b(d)$ to be positive, decreasing and strictly log-concave in b ; and so that $-\pi'_b(d)/\pi_b(d)$ is positive and increasing in d . We also assume



(a) Decomposition by the number of inventors in CPC subclass



(b) Decomposition by the markups of NBER technology sectors

Figure 3: Event study of the application of breakthrough patent on fraction of coauthors on patents who are new connections by concerns of business-stealing

Notes: The figure plots the treatment effect of a breakthrough patent on inventor collaborations using the Borusyak et al. (2021) imputation estimator. The event is the application for a patent by an inventor that is among the 2 percent most cited patents filed in its CPC subclass within a range of 2 years; the event is timed by the year of its application. The outcome is the fraction of new coauthors among all coauthors in a given year of an inventor. Figure 3a shows treatment effect decomposed by above and below median number of inventors in a CPC subclass. Figure 3b shows the treatment effect decomposed by above and below median price-markup in the NBER technology sector of the inventor. Pre-trends are estimated using OLS on not-yet-treated observations. For years with no patent, the value is imputed using the last observed value. Bars indicate the 95% confidence interval. The estimation controls for inventor, experience, and for $\text{year} \times \text{CPC subclass} \times \text{country}$ specific shocks.

that $\pi_b(d)$ tends to 0 as d tends to infinity. Thus, increased numbers of collaborators lead to lower profits, and eventually to full dissipation.

Finally, take $-\pi'_b(d)/\pi_b(d)$ to be increasing in b , so that inventors who see more threat of business stealing see a greater marginal percentage loss in profits from a successful project from adding new collaborators.

We take all functions to be differentiable.

4.2 An Inverted-U Relationship Between Competition and Collaboration

First, let us explain why this simple model lead to an the inverted-U relationship between the potential for business stealing and the number of new collaborators on a project.

The expected payoff to the inventor who has d collaborators is

$$P(d)\pi_b(d).$$

A necessary condition for an optimal d , when it lies strictly between 0 and n (and ignoring integer constraints), is

$$P'(d)/P(d) = -\pi'_b(d)/\pi_b(d). \quad (7)$$

This has a natural interpretation. The left hand side is the elasticity of the probability of a breakthrough and that has to offset the right hand side which is the elasticity of the loss in profits.

Given that $P(\cdot)$ is positive, increasing, and strictly concave, it follows that $P'(d)/P(d)$ is positive and is decreasing in d . Given that $-\pi'_b(d)/\pi_b(d)$ is positive and increasing in d , it follows that there is at most one interior solution to (7). Note also that expected profits are positive at $d = 0$ and converge to 0 as d becomes too large, and so if there is no solution satisfying (7), then the solution is at $d = 0$.

Let $d^*(b)$ denote the unconstrained optimum (ignoring the available n). Since the left hand side of (7) is independent of b and decreasing in d and the right hand side is increasing in b and d , it follows that $d^*(b)$ is decreasing in b (whenever nonzero), as pictured in Figure 4.¹³

¹³The optimum is the intersection of a downward sloping curve that does not depend on b , and an upward




Figure 4: The optimal number of collaborators reflects two effects. Higher levels of potential business stealing reduces the rate at which profits decrease with collaboration, shifting the right hand curve inwards.

This analysis shows the two competing effects. The benefits from new collaborations comes through the elasticity in the probability of a breakthrough, while the fear of business stealing is reflected in the potential lost profits (here the upward shift of the elasticity of the lost profits).

The overall maximum is $\min[n, d^*(b)]$. As an example, if $P(d) = \min\{1, ad\}$ and $\pi_b(d) = \max\{0, \pi_0 - bd\}$ then $d^*(b) = \pi_0/(2b)$ and the overall maximum given the constraints (ignoring integer constraints) is $\min\{n, \pi_0/(2b), 1/a\}$.

Substituting n as one possible proxy for potential business stealing b , the overall maximum becomes $\min[n, d^*(n)]$. If $d^*(n)$ is ever positive, then noting that $d^*(n)$ is decreasing in n , the overall function has an inverted-U shape with a maximum at the point at which $d^*(n) = n$, as it is initially constrained by the min function and eventually follows $d^*(n)$. Figure 5 pictures this for the example above in which the overall optimum is $\min\{n, \pi_0/(2n), 1/a\}$.

Although we have shown this without worrying about the integer constraints on d , the optima for the above functions will be an adjacent integer to the unconstrained solution, and so a similar result will hold, at least in some approximate sense, for most functions satisfying our assumptions.

This establishes an inverted-U shape in the number of added collaborators to a promising project for a very simple example that trades off the probability of a breakthrough against potential lost profits.

sloping curve that shifts upward as b increases, so the equilibrium quantity $d^*(b)$ decreases as b increases, unless it was already 0. This is analogous to shifting up a supply curve intersecting a demand curve that is not shifting.




Figure 5: The optimal number of collaborators d^* as a function of the number of inventors n in the area. For probability of success $P(d) = \min\{1, ad\}$ and profits as a function of the number of collaborators $\pi_n(d) = \max\{0, \pi_0 nd\}$.

4.3 The Increase and Subsequent Decrease in Added Collaborators

Next, we discuss how to extend this model to cover the event study results, incorporating time and information.

There are three periods $\{0, 1, 2\}$, and an inventor is working on a project. There are two states of the world *High*, *Low*, that describe the project's potential. The prior probability of the *High* state is $\theta \in (0, 1)$. The project has the potential to lead to a profit in period 2 only if the state of nature is *High* and not if the state of nature is *Low*.

In period 1 the inventor gets a preliminary signal S_1 about the potential of the project, and the signal takes on one of two values: $\{H, L\}$. If the state of nature is *High*, there is a probability $p \in (0, 1)$ that the signal is $S_1 = H$ and $1 - p$ that the signal $S_1 = L$. If the state is *Low* the probability is 1 that the signal is $S_1 = L$. Thus, conditional on seeing H , the inventor knows the state is *High*, while conditional on seeing L , the posterior that the state is *High* is

$$\theta \frac{1 - p}{1 - \theta p} < \theta. \quad (8)$$

After seeing the signal in period 0, in period 1 the inventor can choose a number of up to n collaborators to help with the project. Invited collaborators observe the signal and decide whether to join the project in period 1.

If the state of nature is *High*, then the project has a probability of $P(d)$ of being

successful in period 2 (as above). If the state of nature is *Low*, the project leads to 0 profits no matter how many people work on it.

If successful, the project yields profits $\pi_b(d)$ where b is the level of potential business stealing.

Take projects without an H signal to be non-breakthrough projects. They do not attract new collaborators. A breakthrough project then attracts new collaborators, and in a way that matches the empirics: the maximal amount for intermediate levels of competition.

The post-phase empirical results – in which fewer new collaborators are added – can then be rationalized, by adding an additional period to the model.

Suppose that in an additional period 3, a successful project can be further developed to result in additional profits by more efforts. So, the innovator can invite a new set of collaborators m to develop the product. The additional profits happen with probability $P^m(d)$ and then generate profits of $\pi_b^m(d)$, that have similar properties as $P(d), \pi_b(d)$, but now in m . However, we presume that the product having already been successful is already known to have very high potential for additional profits, and so the additional gains from adding new team members does less to increase the probability but more to decrease the profits, and so the $m^*(d, b)$ is lower than before for all levels of competition, and moreover is further decreasing in b . Thus, the additional period brings in fewer collaborators than during the first innovation stage.

We note that in this model, we have not distinguished between new and existing collaborators, which could be added with some complications of notation and analysis. We also note that we presume that projects with L signals that are not breakthroughs have no collaboration, which is counter to the data. This is easily accommodated by added a medium signal, which is indicative of a positive but lower profit project that would attract fewer collaborators.

We present a further extension to endogenous competition in the appendix (B).

5 Conclusion

In this paper, we uncovered an inverted-U relationship between the extent to which inventors work with new coauthors and the potential business stealing that face. Second,

we showed that breakthrough innovations lead to fewer new collaborators afterwards, with a bigger drop for bigger threats of business stealing.

While business stealing explains our observations, our event studies do not rule out alternative explanations. One alternative explanation for the basic event study finding is that inventors do not have incentives to share the pie, and consequently, do not expand their coauthorship network once they have a breakthrough. However, this mechanism fails to reconcile the finding that all four of our measures of potential business stealing lead to a bigger decrease in the rate of collaboration.

A second possible explanation is that inventors on breakthrough projects plan multiple patents during their collaboration, and apply for them in the years following the filing of the breakthrough. This could mechanically reduce observed coauthorship in subsequent years, since we capture only those years where inventors do file a patent. However, our findings are that the pattern continues for many years.

A related potential explanation is that inventors might reduce their research efforts and prioritize to consume the monopoly rents earned from breakthrough inventions. Such inventors may choose to collaborate with existing coauthors rather than spend effort in expand their network further. While this explanation may hold for older researchers, subsequent to a breakthrough, younger team members may continue to become team leaders (e.g., Akcigit et al. 2018). Examining patterns of expanding coauthorship among younger researchers would be a natural next step to our analysis.

A third potential explanation is reputation. Having filed a breakthrough patent, inventors may enjoy the privilege of becoming selective about who they work with, and prioritize collaborating with other *high* type inventors. In such a case, we should similarly observe a decline in new coauthorships. High competition areas may have higher gains from reputation due to larger prize pools, and therefore, inventors may set greater thresholds for whom they subsequently collaborate with. This is empirically difficult to disentangle from concerns for business-stealing.

Lastly, our analysis is limited to inventors who have at most one breakthrough patent. It would be interesting to examine how this works for superstar inventors with many breakthroughs and whether there is any selection effect in our data.

Our analysis can also be extended in several other directions. One avenue is to continue to explore the implications of business-stealing for the structure of research networks and the conditions under which successful breakthrough innovators become

central to the network rather than moving to the periphery. Another avenue is to explore policy implications: how and when should collaborative research be encouraged in fields, and how does that depend upon competition? Which competitive environments maximize social welfare when endogenous research is taken into account as well as competition in the product markets? A third avenue is to look in more detail at past coauthorships: when are relationships continued rather than severed? It would also be interesting to directly measure knowledge spillovers and actual business stealing when it does occur, and to see how they interact. Also, although we have focused on patents, there are much wider areas where research is collaborative (across science generally), and where there are some competitive aspects to the process. It would be interesting to see whether similar results hold outside of patents. These and other extensions of the analysis in this paper are left for future research.

References

- Aghion, Philippe, Bloom, Nicholas, Blundell, Richard, et al., 2005. "Competition and Innovation: An Inverted-u Relationship". *The Quarterly Journal of Economics* 120.2, pp. 701–728.
- Aghion, Philippe and Tirole, Jean, 1994. "The management of innovation". *The Quarterly Journal of Economics* 109.4, pp. 1185–1209.
- Akcigit, Ufuk, Caicedo, Santiago, Miguelez, Ernest, et al., 2018. "Dancing with the stars: Innovation through interactions". *NBER Working Paper* 24466.
- Azoulay, Pierre, Graff Zivin, Joshua S, and Wang, Jialan, 2010. "Superstar extinction". *The Quarterly Journal of Economics* 125.2, pp. 549–589.
- Blondel, Vincent D, Guillaume, Jean-Loup, Lambiotte, Renaud, et al., 2008. "Fast unfolding of communities in large networks". *Journal of Statistical Mechanics: Theory and Experiment* 2008.10, P10008.
- Bloom, Nicholas, Schankerman, Mark, and Van Reenen, John, 2013. "Identifying technology spillovers and product market rivalry". *Econometrica* 81.4, pp. 1347–1393.
- Borusyak, Kirill, Jaravel, Xavier, and Spiess, Jann, 2021. "Revisiting event study designs: Robust and efficient estimation". *arXiv preprint arXiv:2108.12419*.
- Burt, Ronald S, 2004. "Structural holes and good ideas". *American Journal of Sociology* 110.2, pp. 349–399.

- Dasaratha, Krishna, 2019. "Innovation and Strategic Network Formation". *arXiv preprint arXiv:1911.06872*.
- De Chaisemartin, Clément and d'Haultfoeuille, Xavier, 2020. "Two-way fixed effects estimators with heterogeneous treatment effects". *American Economic Review* 110.9, pp. 2964–96.
- De Loecker, Jan, Eeckhout, Jan, and Unger, Gabriel, 2020. "The rise of market power and the macroeconomic implications". *The Quarterly Journal of Economics* 135.2, pp. 561–644.
- Dorn, David, Hanson, Gordon H, Pisano, Gary, et al., 2020. "Foreign competition and domestic innovation: Evidence from US patents". *American Economic Review: Insights* 2.3, pp. 357–74.
- Goodman-Bacon, Andrew, 2021. "Difference-in-differences with variation in treatment timing". *Journal of Econometrics* 225.2, pp. 254–277.
- Hong, Lu and Page, Scott E, 2004. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers". *Proceedings of the National Academy of Sciences* 101.46, pp. 16385–16389.
- Immorlica, Nicole, Lucier, Brendan, and Sadler, Evan, 2014. "Sharing Rival Information". *Working paper*.
- Jaravel, Xavier, Petkova, Neviana, and Bell, Alex, 2018. "Team-specific capital and innovation". *American Economic Review* 108.4-5, pp. 1034–73.
- Marco, Alan C, Carley, Michael, Jackson, Steven, et al., 2015. "The USPTO historical patent data files: Two centuries of innovation". *Available at SSRN* 2616724.
- Monath, Nicholas and McCallum, Andrew, 2015. "Discriminative hierarchical coreference for inventor disambiguation". *PatentsView Inventor Disambiguation Technical Workshop, USPTO, Alexandria, VA*. Vol. 122.
- Ozman, Muge, 2009. "Inter-firm networks and innovation: a survey of literature". *Economic of Innovation and New Technology* 18.1, pp. 39–67.
- Page, Scott E, 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press.
- Stein, Jeremy C, 2008. "Conversations among competitors". *American Economic Review* 98.5, pp. 2150–62.
- Walker, Gordon, Kogut, Bruce, and Shan, Weijian, 1997. "Social capital, structural holes and the formation of an industry network". *Organization science* 8.2, pp. 109–125.

A Figures




Figure 6: Distribution of the number of inventors in CPC subclasses

Notes: The figure plots the histogram of the total number of inventors per CPC subclasses using USPTO Patentsview database from 1976 to 2015. The dashed vertical line represents the average number of inventors in CPC subclasses.




Figure 7: Event study of application for a breakthrough patent on team size

Notes: The figure plots the treatment effect of a breakthrough patent on inventor collaborations using the Borusyak et al. (2021) imputation estimator. The event is the application for a patent by an inventor that is among the 2 percent most cited patents filed in its CPC subclass within a range of 2 years; the event is timed by the year of its application. The outcome is the fraction of new coauthors among all coauthors in a given year of an inventor. Pre-trends are estimated using OLS on not-yet-treated observations. For years with more than one patent, the value is the average team size. For years with no patent, the value is imputed using the last observed value. Bars indicate the 95% confidence interval. The estimates control for inventor, experience, and for year×CPC subclass×country specific effects.

B Extending the Model to Have Endogenous Competition

Further insights into business stealing come from extending the model to allow for *endogenous* competition from collaborators, as developed here.

Let us return to the decision of the inventor at time 1 after seeing a signal of type H . Competition now is allowed to depend upon the number of coauthors, and not just the size of the community (in line with Aghion and Tirole 1994).

For the sake of illustration we look at a particular parametric form. Let

$$P(d) = \min\{1, ad\},$$

where $a > 0$, and

$$\pi_b(d) = \max\{1 - b(d, n), 0\}.$$

The threat of business stealing, $b(d, n)$, is itself endogenous and depends on both n and d : some collaborators turn into competitors. In particular:

$$b(d, n) = c_0 + \delta n + \gamma d.$$

The innovator then solves

$$\max_d P(d) \pi_{b(d, n)}(d),$$

or, equivalently,

$$\max_{d \leq \frac{1}{a}, \frac{1-c_0-\delta n}{\gamma}} ad(1 - c_0 - \delta n - \gamma d),$$

presuming that $1 - c_0 - \delta n > 0$ and ignoring integer constraints. This yields:

$$d^* = \min \left[n, \frac{1 - c_0 - \delta n}{2\gamma}, \frac{1}{a} \right].$$

Here again, we see that d^* has an inverted-U shape in n , and is decreasing in the base amount of competition $c_0 + \delta n$ after the cap on n is no longer binding (and for small enough δ).

We also note that the greater the threat of endogenous competition via larger γ , the lower the amount of collaboration that there is. This shows the power of endogenous competition in reducing collaboration.

Again, this can be extended to the post-breakthrough period, by having a higher a , c_0 , δ and γ , which then all work to reduce the size of the collaboration.