


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijioDo patents enable disclosure? Evidence from the invention secrecy act [☆]Gaétan de Rassenfosse ^{a,*}, Gabriele Pellegrino ^b, Emilio Raiteri ^c^a College of Management of Technology, Ecole polytechnique fédérale de Lausanne, Lausanne, Switzerland^b Catholic University of the Sacred Heart, Milan, Italy^c School of Innovation Sciences, Eindhoven University of Technology, Eindhoven, Netherlands

ARTICLE INFO

JEL classification:

O31
O33
O34

Keywords:

Disclosure
Follow-on invention
Knowledge diffusion
Patent

ABSTRACT

This paper provides empirical evidence suggesting that patents may facilitate knowledge disclosure. The analysis exploits the Invention Secrecy Act, which grants the U.S. Commissioner for Patents the right to prevent the disclosure of new inventions that represent a threat to national security. Using a two-level matching approach, we document a negative and large relationship between the enforcement of a secrecy order and follow-on inventions, as captured with patent citations and text-based measures of invention similarity. The effect carries over to after the lift of the secrecy period, suggesting a lost generation of inventions. The results bear implications for innovation and intellectual property policy.

1. Introduction

The patent system is supposed to “promote the progress of science” in two primary ways.¹ Firstly, it enhances the appropriability of inventions. By granting inventors a temporary monopoly on their inventions, the patent system mitigates incentive problems stemming from the public good nature of knowledge (Nordhaus, 1967). Secondly, it facilitates the dissemination of knowledge. The patent document discloses the invention’s technical aspects, with the aim of stimulating follow-on inventions (Ouellette, 2012).

In accordance with the enablement requirement in patent law, inventors must provide a comprehensive description of their invention so that a person skilled in the relevant field can successfully implement it.² This disclosure is essential for delineating the boundaries of the claimed property rights and enables third parties to manufacture and market the invention once the patent has expired. Additionally, third parties may learn from the disclosed information, employing it as a foundation for subsequent inventions. Economists have paid particular attention to this latter aspect, driven by its implications for ‘knowledge spillovers’ (Marshall,

[☆] This paper has benefited from comments from participants at the EPIP conference at the University of Bordeaux (2017), the InnoPat conference at ZEW (2019), workshops at the University of Luxembourg (2018), the University of Bath (2018), the University of Strasbourg (2018), as well as seminars at the University of Maastricht (2019), the Complutense University of Madrid (2019), the University of Barcelona (2019), the Max Planck Institute for Innovation and Competition (2019), Utrecht University (2020) and the NBER Productivity Seminar (2020).

* Corresponding author.

E-mail address: gaetan.derassenfosse@epfl.ch (G. de Rassenfosse).

¹ U.S. Constitution art. I, §8, cl. 8.

² U.S. Code Title 35 §112.

<https://doi.org/10.1016/j.ijindorg.2023.103044>

Received 28 December 2021; Received in revised form 11 December 2023; Accepted 14 December 2023

Available online 20 December 2023

0167-7187/© 2023 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

1920; Krugman, 1992; Jaffe et al., 1993). Nonetheless, as Sampat (2018, p. 20) noted in his comprehensive literature review, “the quantitative effects of the disclosure function of patents on rates of innovation are not well known.”

Following a body of scholarship (Rantanen, 2012; Lemley, 2012; Hall et al., 2014), one can conceptualize the ‘disclosure effect’ of patents in two distinct ways. In a narrow way, the disclosure effect pertains to unveiling the technical information in the patent document. This viewpoint aligns with the conventional understanding of disclosure. In a broader way, the disclosure effect encompasses not only the conventional aspect but also recognizes that patent protection facilitates the revelation of *additional* information about the invention, termed ‘peripheral’ disclosure by Rantanen (2012). Patents guard inventors against the threat of expropriation. Consequently, once an inventor obtains a patent, she no longer needs to keep the invention secret. She can more freely disclose details of the invention, e.g., in trade exhibitions or scientific publications, while retaining the ability to prevent others from utilizing the invention.

The central aim of this paper is to quantify the extent of the disclosure function of patents in its broader sense, compared to the baseline of maintaining secrecy. Our approach leverages the Invention Secrecy Act of 1951, which empowers the U.S. Commissioner of Patents with the authority to withhold the publication of patent documents posing a threat to national security. Inventors of patents subject to a secrecy order are legally prohibited from divulging any material information relevant to the invention. In essence, all avenues for disseminating information are stifled—the invention disappears into a figurative ‘black hole’ for the duration of the secrecy period. We have identified a total of 2,542 patents that were subjected to a secrecy order during the period spanning from 1982 to 2000. Subsequently, these secrecy orders were rescinded by the Commissioner of Patents, allowing us to investigate the impact of disclosure (or, more precisely, the absence of it) on follow-on inventions.

The empirical analysis in this study relies on a matching approach, estimating regression models on a carefully selected subsample of control group observations. Within this framework, treated patents are patents cited by a patent subject to a secrecy order. Control patents, on the other hand, are patents that are matched to the treated patents. The main outcome variable is the number of citations received by the treated and control patents, as a proxy for follow-on inventions. We also measure follow-on inventions using a text-based measure of invention similarity.

Our findings reveal a robust and significant negative relationship between the enforcement of a secrecy order and the emergence of follow-on inventions. The estimates suggest that the imposition of a secrecy order lowers the arrival of follow-on inventions by 40 percent during the period of the secrecy order, relative to the baseline of no secrecy orders. Remarkably, this adverse effect persists even after the lift of the secrecy order, suggesting a ‘lost generation’ of follow-on inventions. Our results are consistent when employing the text-based measure of invention similarity.

We validate the findings through a battery of robustness tests, including, but not limited to, alternative regression models, a falsification test, stringent sample restrictions, and matching on technological trajectories. To address potential concerns that secrecy orders may primarily affect the patenting behavior of firms in the defense industry rather than fundamentally curbing inventions, we investigate follow-on inventions by non-defense firms. We observe a similar decline in patenting activity among non-defense firms (defined in various ways). These results, along with other findings presented in this paper, collectively provide evidence of a genuine reduction in follow-on inventions after enforcing a secrecy order.

Given the pivotal role of innovation in driving economic growth, our empirical findings hold considerable significance for public policy. The debate surrounding whether disclosure within the patent system facilitates subsequent inventions has been a prominent subject of discussion within the scholarly literature. This paper contributes original empirical evidence that supports the affirmative stance, indicating that, indeed, disclosure fosters follow-on inventions. Furthermore, our findings spotlight one of the social costs associated with secrecy: a reduction in the generation of follow-on inventions.

The remainder of this paper is structured as follows. Section 2 furnishes background information on the disclosure function of the patent system. Section 3 elaborates on the identification strategy employed, while Section 4 provides a comprehensive overview of the data and methods utilized in our analysis. Section 5 delves into a detailed discussion of the econometric results. Finally, Section 6 offers concluding remarks.

2. Literature background

A recent strand of research has contributed empirical insights into the disclosure function of patents through the use of quasi-experimental designs. Furman et al. (2021) leverage the openings of regional patent libraries by the United States Patent and Trademark Office (USPTO) between 1975 and 1997. They document a noteworthy upsurge in local patenting activities subsequent to the establishment of these libraries, attributing this phenomenon to the enhanced disclosure of technical knowledge facilitated by these libraries. Cox (2019) observes a similar effect with the introduction of the disclosure requirement in British patent law in 1734. Baruffaldi and Simeth (2020), as well as Hegde et al. (2020), exploit the introduction of the American Inventors Protection Act (AIPA), which significantly reduced the publication time for patent documents. Both studies report substantial increases in the flow of information from patents following the enactment of the AIPA, signifying noteworthy dissemination of technology.

These aforementioned studies yield evidence of a positive impact of ‘conventional’ disclosure on the generation of follow-on inventions. They gauge the effects of enhancing access to patent literature against a baseline in which access to patent documents is either costly or temporarily unavailable, but the legal consequences of patents remain in effect. This finding is particularly remarkable, for it suggests that inventors glean valuable insights from the patent literature.

In contrast to previous studies, inventions subject to a secrecy order were expected to receive patent protection and be disclosed via the patent system. However, the award of patents was delayed, and the peripheral disclosure channels were silenced. A study by Gross (2022) offers an analysis closely aligned with ours. The author exploits a mass patent secrecy program that persisted

throughout World War II (WWII). Gross's analysis sheds light on the impact of secrecy at the *intensive margin*, demonstrating that patents subjected to *longer secrecy terms* were less likely to be cited (and used) by future patents. Nevertheless, as Gross explains, it remains unclear "why a temporary invention secrecy policy in the 1940s would affect long-run, post-war citations, and especially citations from a population of patents which were mostly filed after secrecy orders were rescinded." Indeed, due to data limitations, the paper does not address the flow (or lack thereof) of knowledge during the enforcement of secrecy orders or immediately thereafter. Additionally, the context of WWII was marked by substantial investments and rapid advancements in military technologies, leaving open questions about the generalizability of these results outside the context of a global war. Another notable departure from Gross (2022) is that we offer an analysis at the *extensive margin*. We gauge the impact of secrecy orders against the baseline of disclosure, rather than comparing it to shorter periods of secrecy.

3. Definition of 'secret inventions' and 'follow-on inventions'

We aim to assess the extent to which knowledge flows are disrupted when an invention is kept secret rather than patented. In an ideal experiment, we would observe the same invention in two mutually exclusive regimes: one in which the inventor files a patent application for the invention and another in which the inventor opts to maintain the same invention as a secret. We would then compare how many inventions were built upon the focal invention in both regimes. Inventions 'building' upon the secret invention would occur because knowledge is 'in the air,' and the difference with the non-secret invention would be the broad disclosure effect. While it is infeasible to conduct this ideal experiment, our discussion highlights the necessity of two main components: a means of identifying secret inventions and a method for tracking follow-on inventions.

To identify secret inventions, we leverage the 1951 Invention Secrecy Act, which remains in effect to this day.³ This Act allows the patent office to impose secrecy orders on patent applications. These secrecy orders also oblige inventors to refrain from disclosing the invention, under the threat of a potential jail sentence of up to two years. Consequently, the secrecy regime for such inventions shares specific characteristics with inventions kept as trade secrets, at least from the perspective of information disclosure. However, a critical distinction exists: secrecy orders have a finite duration. They last for one year and can be extended for as many years as deemed necessary by the relevant government agency, until the invention is no longer a threat to public safety and defense. At that point, the patent application is approved, and the invention is disclosed.

As for tracking subsequent inventions, we draw upon the well-established tradition in the economic literature that utilizes patent citations as "paper trails" of knowledge diffusion (Jaffe et al., 1993). Despite some criticisms (Alcacer and Gittelman, 2006; Kuhn et al., 2020), citation data remain the most comprehensive and frequently utilized metric for monitoring cumulative innovation in economic research (e.g., Belenzon, 2012; Galasso and Schankerman, 2014; Moser et al., 2017; Jaffe and de Rassenfosse, 2017). We follow this tradition and employ the count of forward citations as our primary measure for tracking follow-on inventions. Additionally, we will employ an alternative approach involving a content-based measure. This approach utilizes data provided by the Google Patents team, assessing the textual similarity of patents. This measure derives from a model that has learned a set-of-words embedding of the full text of a patent relative to its technology class (CPC) using the WSABIE embedding algorithm (Weston et al., 2011).⁴ From a technical viewpoint, this method is superior to other approaches that have used unsupervised embeddings relying either on single-layer Neural Networks such as Word2Vec (Whalen et al., 2020) or more basic methods such as one-hot encoding (Iaria et al., 2018; Arts et al., 2018) and TD-IDF (Younge and Kuhn, 2016). See de Rassenfosse and Raiteri (2022) for previous use of Google embeddings.

While both components are in place, we still face a major challenge. If an invention is kept secret, there are no forward citations to count, at least not until the lift of the secrecy order. Knowledge generation is a cumulative process (e.g., Scotchmer, 1991; Jones, 2009; Dosi and Nelson, 2010). By preventing a specific body of new knowledge from emerging, we not only deprive the world of that piece of knowledge but also of potential improvements and modifications that could build upon that knowledge. Consequently, we take a step back in the invention sequence and identify the technological antecedents to secret inventions.⁵ We can view these technological antecedents as the starting point of an invention sequence that can evolve in two alternative regimes: the *secret* regime and the *patent* regime. In the *secret* regime, the technological antecedents are cited (*i.e.*, used as input) by an invention that receives a secrecy order, resulting in no patent issuance and no disclosure. In the *patent* regime, the subsequent invention is patented and disclosed. In both regimes, the invention sequence that commenced with the antecedents continues to evolve over time. Even in the *secret* regime, we can observe new follow-on inventions that cite the antecedents during the period between the imposition of the secrecy order on the follow-on patent and the lift of the secrecy order.

Our empirical strategy lies in estimating the differences in the rate of emergence of follow-on inventions in this time window between the two regimes, as summarized in Fig. 1.

³ See Manual of Patent Examining Procedure, 120.

⁴ See, e.g., <https://patents.google.com/?q=~patent%2fUS7945525B2>. More details on the similarity algorithm are available at <https://media.epo.org/play/gsgoogle2017>.

⁵ Secret inventions are examined before Notices of Allowability are issued. Patent examiners produced a search report with the citations to the relevant prior art as of the time of examination. We can thus use backward citations to track the technological antecedents of secret patents in the inventive sequence.

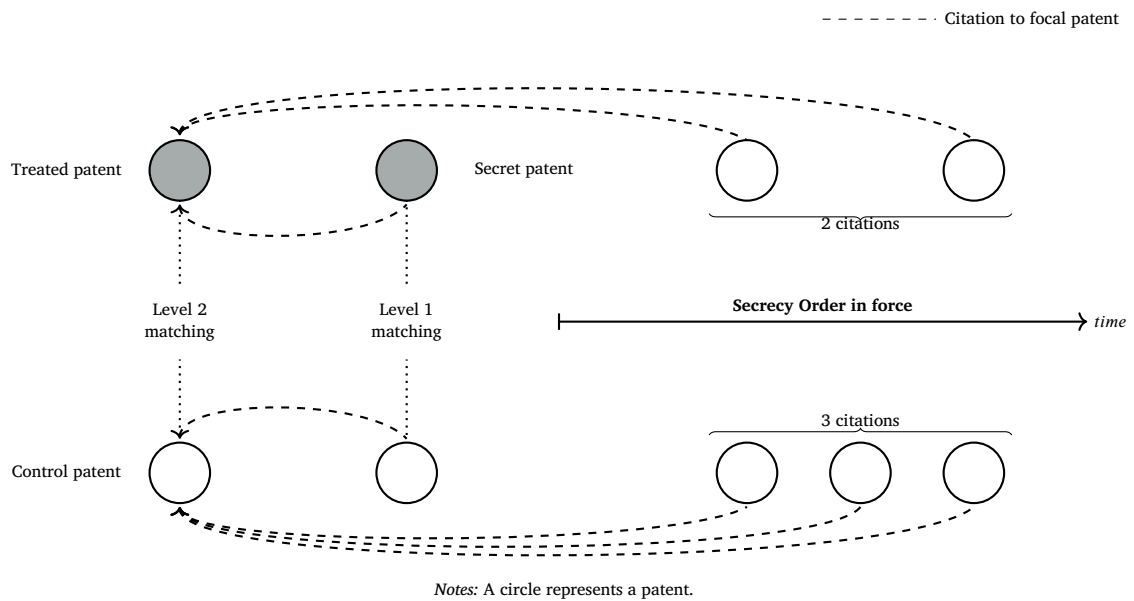


Fig. 1. Illustration of empirical strategy.

4. Data

We leverage the USPTO Patent Examination Research Dataset (PatEx) to identify patents that were initially subject to a secrecy order but subsequently had the order rescinded. The PatEx dataset, compiled by Graham et al. (2018), provides us full access to the extensive data collected by the Public Patent Application Information Retrieval system (Public PAIR).⁶ The PAIR system offers comprehensive data on the administrative transactions that transpired between patent examiners and applicants throughout the patent prosecution process, including information on secrecy orders.

Our sample encompasses 2,542 patents filed between 1982 and 2000 that were subjected to a secrecy order. Commencing with patents filed in 1982, we have the ability to trace the communications used by examiners to notify applicants of the imposition of a secrecy order, including renewals and the subsequent rescission of secrecy orders.⁷ Fig. A.1–A.3 in Appendix A provides an example of such communications. To maintain the integrity of our analysis and avoid potential biases associated with the American Inventors Protection Act (AIPA), we limit our sample to patents filed up to the year 2000.⁸ Of these patents, 421 were associated with foreign assignees, which we exclude to strengthen the identification strategy, resulting in a sample size of 2,121 patents.

As detailed in Section 3, we construct our estimation sample through a two-tiered matching process, illustrated in Fig. 1. In essence, we are carefully trimming observations to construct a meaningful control group, as discussed *inter alia* by Ho et al. (2007, 2011) and Iacus et al. (2012, 2019). The first tier, termed *Level 1* matching, involves identifying an appropriate group of control patents for the sample of 2,121 patents subject to a secrecy order. We draw these control patents from a pool of 123,618 U.S. patents filed between 1982 and 2000, all of which were examined by patent examiners who had inspected at least one patent subject to a secrecy order. This approach ensures some consistency in the subject matter of the inventions. Subsequently, we implement an exact matching procedure based on the primary technological class (USPC) and the year of application. This procedure results in a sample of 46,291 *Level 1* control patents matched with 1,941 patents subject to a secrecy order.⁹

In the second tier, termed *Level 2* matching, we begin by identifying all patent antecedents associated with the two selected groups, comprising secret patents and their matches. This process yields 16,146 patent antecedents linked to the group of secret patents and 410,377 patent antecedents linked to the group of matched patents. Subsequently, we employ an exact matching technique to enhance homogeneity between these two groups of patents. We perform exact matches based on the primary USPC patent class, the year of patent application, and the geographical location of assignees (U.S. vs. foreign). Furthermore, we impose that both treated and control patent antecedents must have received an identical number of forward citations from the year of their application until

⁶ The PatEx dataset is available at <https://www.uspto.gov/learning-and-resources/electronic-data-products/patent-examination-research-dataset-public-pair>.

⁷ Information on the rescission of the secrecy order is not available prior to 1981; thus, patents might have been subject to a secrecy order that was then lifted without our knowledge.

⁸ After the AIPA, the patent application is disclosed in full after 18 months from filing. See Graham and Hegde (2015) for a detailed description of this policy change.

⁹ We do not find any exact match for 180 treated patents.




Fig. 2. Frequency distribution of secrecy order length and secrecy order's application year.

the year in which the secrecy order was imposed on the citing patent.¹⁰ This restriction ensures that we consider treated and control patents that appear similarly ‘promising’ until the imposition of the secrecy order. By design, the treated and control antecedents are cited at least once, implying that we focus our attention on inventions that offer some proven potential for follow-on inventions. Our method is silent on inventions disclosed but never built upon, reminding us that disclosure does not necessarily imply the arrival of follow-on inventions (e.g., some inventions might be technological dead-ends or be of limited value).

This stringent double-matching procedure yields a final sample of 1,313 treated and 3,176 control patent antecedents, which constitute our focal patents. These patents are linked to 725 citing patents with a secrecy order and 2,824 citing patents without a secrecy order, respectively. Additionally, we enrich our dataset by retrieving relevant characteristics of both our focal patents and citing patents, including family size and the number of claims, backward citations, and inventors listed in the patent document (sourced from the PATSTAT database, release 2015).

Fig. 2a presents the distribution of secrecy order durations (solid line). This distribution is right-skewed, with a mode of one year and an average duration of 4.1 years. Fig. 2b plots the frequency distribution of patent filing years for the 725 patents subjected to secrecy orders. The majority of patents in our sample originate from the late 1980s to the early 1990s.

To provide insight into the technological domains of these patents, Table B.1 in Appendix B presents different USPC main classes sorted by the number of patents. Notably, many of the most frequent classes pertain to technologies with obvious military applications, such as class 342 (Communications: directive radio wave systems and devices), class 102 (Ammunition and explosives), class 149 (Explosive and thermic compositions or charges), and class 343 (Communication: radio wave antennas). However, this is not true for all technologies. For instance, a declassified document from 1971 reveals that technologies such as “solar photovoltaic generators,” “computer-aided design,” and the “rapid production of photographic prints” were considered for restrictions.¹¹

Table 1 Panel A provides summary statistics on citing patents with secrecy orders for the total sample (2,121 observations) and the matched sample (725 observations). Focusing on the matched sample, we observe that, on average, each patent contains approximately 15 claims, has received over 16 citations, and has a family size of two, signifying that it has been extended to other jurisdictions following the rescission of the secrecy order. Moreover, each patent references about twelve prior patent documents and one non-patent literature document, including scientific references. A team of two inventors typically produces these patents. The average duration of secrecy orders is 3.7 years; however, this variable exhibits a right-skewed distribution, ranging from several months to 20 years, with over one-third of secrecy orders lasting less than two years (see Fig. 2a, dashed line). In this latter case, it is clear that the secrecy order did *not* extend the non-disclosure period because patent documents were not published before grant until the AIPA came into force in November 2000. In fact, the median lag between the filing date of a patent application at the USPTO and the grant date is about 23 months in the period we consider. Therefore, for our baseline empirical analysis, we focus exclusively on patents with secrecy orders that exceed three years. This criterion ensures that the disclosure of information about the secret patent was substantially delayed compared to the control patent. Consequently, our final estimation sample comprises 505 treated and 1,244 control patent antecedents (our focal patents), linked to 297 citing patents with secrecy orders and 1,152 citing patents without secrecy orders, respectively.

¹⁰ For the calculation of the forward citations of the control patent antecedents, we consider the year of the imposition of the secrecy order of the associated treated patent.

¹¹ Source: <https://fas.org/sgp/othergov/invention/psrcl.pdf>, last accessed Jan 4, 2023.

Table 1
Summary Statistics.

Panel A. Patents with a secrecy order imposed								
	Full sample				Matched sample			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Secrecy length	4.1	4.0	0	23	3.7	3.6	0	20
No. of claims	14.8	12.4	1	126	15.4	11.9	1	106
Family size	2.1	2.5	1	29	2.0	2.3	1	29
No. of backward citations	9.6	8.0	1	100	12.1	9.4	1	100
No. of non-patent literature	1.4	3.4	0	37	1.3	3.1	0	37
No. of inventors	2.0	1.2	1	12	2.0	1.3	1	11
No. of forward cit (tot)	16.5	20.9	0	234	16.6	19.4	0	179
Number of observations	2,121				725			
Panel B. Treated and control focal patents (estimation sample)								
	Treated		Control		Mean-Diff			
<i>Number of citations accrued:</i>								
Pre-Secrecy	2.81	(2.88)	2.81	(2.88)	0.00	[1.000]		
Secrecy	2.33	(3.51)	3.72	(5.96)	-1.39	[0.000]		
Secrecy+3	3.21	(5.26)	5.28	(8.39)	-2.07	[0.000]		
Secrecy+5	3.67	(6.14)	6.23	(9.83)	-2.56	[0.000]		
Priority	0.53	(0.49)	0.53	(0.50)	0.00	[0.865]		
Family size	3.10	(3.31)	3.33	(3.58)	-0.23	[0.286]		
No. of backward citations	6.13	(5.03)	6.27	(6.49)	-0.14	[0.672]		
No. of non-patent literature	0.63	(1.72)	0.72	(2.39)	-0.09	[0.600]		
No. of inventors	1.80	(1.09)	1.85	(1.21)	-0.05	[0.426]		
No. of claims	9.25	(10.19)	9.92	(9.91)	-0.67	[0.203]		
Age classes (USPC)	129.60	(23.57)	129.50	(23.42)	0.09	[0.941]		
Number of observations	505		1,244		1,749			

Notes - Panel B: Standard deviation in parentheses, p-value in square brackets.

5. Econometric results

5.1. Empirical model

Our econometric model is:

$$y_i = c + \phi \text{Secret}_i + \gamma \mathbf{X}_i + \delta_s + \delta_t + \epsilon_i, \quad (1)$$

where the dependent variable y_i represents the number of forward citations received by focal patent i starting from the year in which the secrecy order of its citing patent was imposed. We measure citations at two points: until the publication year of the secret patent (which roughly corresponds to the year of the rescission of the secrecy order) and until five years after the publication of the patent.¹² More details regarding the text-based measure of similarity are discussed in Section 5.3.

The variable of interest is the dummy variable Secret_i . It takes the value 1 if the focal patent has been cited by a patent subject to a secrecy order, and 0 otherwise. The sign and significance of the coefficient ϕ will provide insights into whether and to what extent the enforcement of a secrecy order acts as a barrier to the diffusion of knowledge.

The vector \mathbf{X}_i captures features of the focal patent, including the size of the patent family at the international level, the number of inventors listed on the patent document, the number of claims, and the number of backward citations to both patent and non-patent literature.¹³ The vector δ_s consists of a complete set of USPC (main) class fixed effects, and the vector δ_t includes filing year fixed effects for both citing and cited patents. The parameters c and (ϕ, γ) are the intercept and slope coefficients, respectively. The error term ϵ_i represents unobserved factors and is assumed i.i.d.¹⁴

¹² For the calculation of the forward citations of a control patent, we consider the year in which the secrecy is rescinded for the associated treated patent.

¹³ It is worth noting that a specification including matched-pair fixed effects produces very similar results.

¹⁴ Abadie and Spiess (2022) advise clustering the standard errors at the level of the matched sets in case of potential misspecification of the linear regression model. All our results are robust to matched set clustering.

Table 2
Effect of secrecy orders on forward citations: baseline estimates.

	Secr. length 3 y		Secr. length 6 y		Neg. Binomial	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>No. of fwd citations:</i>	Secrecy	Secrecy+5	Secrecy	Secrecy+5	Secrecy	Secrecy+5
Secret	-1.354*** (0.241)	-2.480*** (0.381)	-1.769*** (0.353)	-2.989*** (0.541)	-0.953*** (0.143)	-1.829*** (0.226)
Pre-Secrecy	0.369*** (0.124)	0.551*** (0.160)	0.431** (0.170)	0.511** (0.213)	0.161*** (0.034)	0.268*** (0.052)
Priority	0.240 (0.389)	0.932 (0.611)	0.208 (0.461)	0.917 (0.829)	0.218 (0.157)	0.627** (0.249)
Family size	0.017 (0.043)	0.079 (0.065)	-0.016 (0.072)	0.063 (0.114)	0.021 (0.021)	0.056* (0.033)
No. of backward citations	0.080* (0.044)	0.124* (0.073)	0.127* (0.076)	0.197 (0.123)	0.018 (0.011)	0.021 (0.018)
No. of non-patent literature	0.055 (0.084)	0.143 (0.143)	0.100 (0.132)	0.156 (0.221)	0.047 (0.038)	0.090 (0.059)
No. of inventors	-0.158 (0.109)	-0.156 (0.176)	-0.193 (0.159)	-0.042 (0.258)	-0.055 (0.051)	-0.044 (0.077)
No. of claims	0.021 (0.013)	0.044* (0.025)	0.015 (0.018)	0.022 (0.036)	0.017** (0.007)	0.040*** (0.011)
Class fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Filing year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Filing year (secr.) fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	1,749	1,749	848	848	1,749	1,749
R^2	0.32	0.33	0.53	0.51	-	-
Mean of dep. var.	3.313	5.479	4.140	6.288	3.313	5.479

Notes: Estimation method is OLS unless otherwise specified. Robust Standard errors in parentheses. Significant at *** 1%, ** 5% and * 10%. The dependent variable records the cumulative count of forward citations that the focal patent i has received from the year in which the secrecy order of its citing patent has been imposed until: the publication year of the patent (cols. 1, 3, and 5); five years after the publication of the patent (cols. 2, 4, and 6). Self-citations and foreign (not-U.S.) citations are excluded from the count of the dependent variables. The variable Secret is a binary indicator that identifies those focal patents that have been cited by at least one patent with a secrecy order. Secrecy length of at least three years in columns (5)–(6). Coefficients in columns (5)–(6) are expressed as marginal effects.

Table 1 Panel B presents summary statistics for control and treated patents in the estimation sample. Notably, treated patents appear to receive fewer citations during the secrecy order period (represented by the variable ‘Secrecy,’ which spans from the imposition of the secrecy order to the publication of the patent document), providing initial evidence of the disclosure effect. Importantly, treated and control patents have the same number of citations until the secret patent arrives, as designed (variable ‘Pre-Secrecy,’ which spans from the focal patent’s filing date to the secret patent’s arrival). The exact *Level 2* matching effectively balances treated and control patents concerning all relevant covariates.

5.2. Baseline results

Table 2 presents the results of the baseline OLS estimates (based on the matched subsample of data) for the main dependent variables. As discussed in Section 4, the sample comprises treated and control focal patents (the ‘antecedents’) that were cited by a patent with a secrecy order (or its control) lasting at least three years. The results, reported in columns (1) and (2), reveal a negative and statistically significant relationship between the enforcement of a secrecy order and follow-on inventions. Specifically, a patent cited by a patented invention that has been kept secret for at least three years receives, on average, approximately 1.35 fewer citations than its control counterpart (column 1). When compared with the average of 3.31 forward citations received by focal patents, this coefficient implies a 41-percent reduction in forward citations during the secrecy period (1.354/3.313). This negative effect of secrecy slightly increases to 45 percent when extending the citation window to five years after the publication of the citing patent, *i.e.*, after the lift of the secrecy order (column 2). Consequently, it appears that lifting the secrecy order does not reverse the trend of lost citations—the decrease persists even after the secrecy order is rescinded.

To gain further insights into the effect of secrecy at the intensive margin, we re-estimate the regression model by considering a secrecy length threshold of at least six years, rather than three. The OLS estimates in columns (3) and (4) of Table 2 result in coefficients of -1.77 and -2.99, respectively, corresponding to a reduction of 43 percent and 48 percent in the number of forward citations received by the focal patent compared to the control patent.

Next, to account for the count nature of citations, we re-estimate the baseline model using the negative binomial estimator, as shown in columns (5) and (6) of Table 2. When considering citations received during the secrecy period (column 5), the estimate implies a 29-percent reduction in the number of forward citations received by the focal patent (and 33% in column 6). We obtain qualitatively similar results with OLS by transforming the dependent variable to the log number of citations (not reported).

Collectively, these initial estimates suggest a substantial effect of secrecy orders on follow-on inventions. Being cited by a patent temporarily kept secret is associated with a significant reduction in the arrival of follow-on inventions. Furthermore, the length of the secrecy order appears to amplify this negative effect.

5.3. Tracking follow-on inventions using data on patent text similarity

As discussed in Section 3, the use of patent citations to track follow-on inventions has faced criticism. This section establishes the robustness of the results when identifying follow-on inventions using a content-based approach.

We begin with the set of treated and control patents identified through the *Level 2* matching method outlined in Section 4. However, instead of relying on citation counts, we track inventions most similar to the treated and control focal patents based on the content of the patent documents. The search for similar patents on the Google Patents platform (see footnote 4) yielded an average of 1,905 patents per document, but we retain information on the first hundred most similar patents.¹⁵ To measure the degree of similarity between the focal patents and these patents, we exploit the embedding vectors available from the Google Big Query data platform. These vectors allow us to calculate the cosine similarity between any two patents, providing a precise ranking of patent similarity.

Before using these data, we filter out certain patents from the list returned by Google Patents. This broad list comprises worldwide patents filed across various time periods and technology fields and includes patents related to the same family as the focal patents. For each focal patent i , we restrict the set to U.S. patents, keep only one similar patent per INPADOC family, and exclude similar patents related to the focal patent i 's family. Furthermore, we narrow down the sample by excluding similar patents with a primary USPC class different from the USPC class of the matched (secret or control) patents citing focal patent i . This refinement helps us better isolate the potential hindering effect caused by the secrecy order on follow-on inventions.

The subsequent step involves constructing two dependent variables, following the same methodology as in the primary analysis. We count patents similar to focal patent i that emerged between the year in which the secrecy order was imposed on its citing patent and the publication year of the secret patent (Secrecy), as well as five years after the patent's publication (Secrecy+5).

In an initial analysis, we restrict similar patents to follow-on patents with a cosine similarity score greater than or equal to 0.80, which corresponds to the median similarity value in the data. The results in Panel A of Table 3, columns (1)–(2), reaffirm the negative effect of imposing a secrecy order on follow-on inventions. Patents cited by patents subject to a secrecy order are associated with significantly fewer textually similar patents compared to the control group. Furthermore, these findings remain robust when focusing on patents with a cosine similarity score greater than 0.85, which corresponds to the 90th percentile value in the data (columns 3 and 4).

5.4. Robustness tests

Several potential sources of bias in our estimates warrant consideration, and this section explores these concerns. One potential threat pertains to the specific nature of patents subjected to secrecy orders. While our empirical analysis centers on patents cited by secret patents (along with their control patents), there may be lingering concerns about whether our two-level matching approach fully mitigates this potential bias. A second concern revolves around the possibility of alternative explanations accounting for the observed patterns. Specifically, technology domains subject to secrecy orders could be well-known among industry insiders, and the reduction in citations might reflect a decline in the number of patent filings rather than a genuine decrease in inventive activity.

5.4.1. Falsification test and alternative matching

While our two-level matching approach is designed to account for the specific characteristics of secret patents, it is important to consider the possibility that certain attributes of a patent could make it more likely to be cited by a secret patent.

To assess the validity of our identification strategy, we can leverage a feature of the data. As illustrated in Fig. 2a (dashed line), some secrecy orders are relatively short, with approximately 100 patents in the matched sample subjected to secrecy orders lasting less than one year. For such short secrecy orders, the impact of secrecy should be null since patents were not published before grant during our study's time frame.¹⁶

We conduct a falsification test by replicating the analysis on the subset of patents (along with their controls) that were subject to secrecy orders lasting less than one year. As shown in Table 3 Panel B, the treatment effect becomes statistically indistinguishable from zero, providing additional support for the validity of our approach (columns 1 and 2).

Furthermore, our results remain robust to a more stringent alternative matching approach. In our baseline analysis, the *Level 1* matching coupled secret patents with similar non-secret patents based on their primary technology class and filing year. However, it

¹⁵ Google Patents only deems two patent documents similar if their cosine similarity score is 0.7 or higher. Hence, our analysis focuses on patents exceeding the 0.7 similarity threshold.

¹⁶ During the time of our study, patents were not published until they were granted, and examination took at least as long as these short secrecy orders. In other words, the secrecy constraint imposed by secrecy orders was not binding.

Table 3
Effect of secrecy orders on follow-on inventions: similarity measures, placebo test, and alternative matching method.

Panel A. Text-based similarity measure				
	Sim. > 0.8		Sim. > 0.85	
	(1)	(2)	(3)	(4)
<i>No. of similar patents:</i>	Secrecy	Secrecy+5	Secrecy	Secrecy+5
Secret	-0.823** (0.321)	-2.247*** (0.666)	-0.502*** (0.182)	-1.123*** (0.356)
Number of observations	1,748	1,748	1,748	1,748
Mean of dep. var.	3.361	8.074	0.956	2.170
Panel B. Placebo and alternative matching				
	Placebo		Alternative Matching	
	Secrecy	Secrecy+5	Secrecy	Secrecy+5
Secret	-0.168 (0.146)	0.236 (0.570)	-1.588*** (0.391)	-2.797*** (0.618)
Number of observations	689	689	534	534
Mean of dep. var.	0.929	4.254	3.160	5.105

Notes: The estimation method is OLS. Robust Standard errors in parentheses. Significant at *** 1%, ** 5% and * 10%. This table only reports the regressor of interest. Regressors not listed are those of Table 2. See main text for additional details.

did not consider the similarity of the knowledge upon which they were built. In this alternative matching, we impose in *Level 1* that a non-secret patent must share at least one backward citation with a secret patent. Subsequently, we identify treated and control focal patents (*Level 2*) from this more restrictive sample using the same covariates as in the baseline case: application year, primary USPC class, and the number of forward citations accumulated from the patent's filing year until the filing year of the associated citing patent application. This process yields a final sample of 534 treated and control cited patents.¹⁷

The coefficients of the variable of interest, presented in Panel B of Table 3, remain negative and statistically significant. Compared to the baseline model, the effect is even more pronounced. The OLS estimate in column (3) suggests a 50-percent reduction in forward citations during the secrecy period and a 55-percent reduction up to five years after the secrecy period concludes (column 4).

5.4.2. Inventions or patents?

As previously alluded to, the observed decline in citations may potentially signify a reduction in patenting rather than a decline in actual inventions. It is plausible that follow-on inventive activity remains strong, but assignees operating in sensitive technological domains may opt for secrecy over patenting (or be forced into secrecy), resulting in a reduction in citations. While we cannot directly address this issue (as we do not observe inventions that were not filed), we can identify assignees who were more likely to be aware of and exposed to secrecy orders or may better understand the sensitive nature of their technologies.

In our first robustness test, we focus on the population of assignees engaged in sensitive technologies, specifically contractors for the Department of Defense (DoD). We partition the dependent variable into two variables. The variable 'Cites by DoD contractors' measures forward citations originating from patents held by contractors that conducted R&D work with the U.S. government between 1982 and 2000.¹⁸ The variable 'Cites by non-DoD contractors' represents the complementary set. Table 4 Panel A displays the results of these additional estimates. Interestingly, the treatment effect appears more pronounced when focusing on citations from organizations that are not DoD contractors compared to those that are DoD contractors. This finding suggests that there is indeed a genuine reduction in follow-on inventions.

Next, we filter by relevant technologies rather than specific organizations. We segment citations into two categories: those stemming from USPC classes susceptible to receiving a secrecy order and those that are not at risk. The variable 'Cites from secret classes' captures forward citations from patents belonging to USPC classes that were assigned as the primary class to patents receiving a secrecy order during our reference period. The variable 'Cites from non-secret classes' exclusively includes forward citations from

¹⁷ We conducted several additional analyses (not reported) to enhance the *Level 2* matching. Specifically, we exploited the backward citation network and the characteristics of the pre-secrecy citations of the treated and control patents to increase the similarity of their technological trajectories. These additional checks consistently led to a more pronounced impact of secrecy as we enhanced the similarity between the technological trajectories of the treated and control patents.

¹⁸ To identify contractors that performed R&D work for the Department of Defense during the period 1982–2000, we utilized the *Records of Prime Contracts Awarded by the Military Services and Agencies*, available in electronic format from the National Archives and Records Administration (NARA). These records contain data from the Defense Contract Action Data System (DCADS) regarding contracts for goods and services between private sector entities and Department of Defense agencies spanning from 1976 to 2000. The data are available at <https://www.archives.gov/research/electronic-records/reference-report/federal-contracts>.

Table 4
Effect of secrecy orders on forward citations: additional robustness checks.

Panel A. Citations from U.S. government contractors or not				
	Cites by DoD contractors		Cites by non-DoD contractors	
	(1)	(2)	(3)	(4)
<i>No. of fwd citations:</i>	Secrecy	Secrecy+5	Secrecy	Secrecy+5
Secret	-0.604*** (0.186)	-1.152*** (0.272)	-0.900*** (0.150)	-1.509*** (0.227)
Number of observations	1,749	1,749	1,749	1,749
Mean of dep. var.	2.197	3.402	1.257	2.309
Panel B. Citations by secret and non-secret classes				
	Cites from secret classes		Cites from non-secret classes	
	Secrecy	Secrecy+5	Secrecy	Secrecy+5
<i>No. of fwd citations:</i>	Secrecy	Secrecy+5	Secrecy	Secrecy+5
Secret	-1.124*** (0.207)	-1.936*** (0.303)	-0.238** (0.112)	-0.554*** (0.200)
Number of observations	1,749	1,749	1,749	1,749
Mean of dep. var.	2.666	4.250	0.632	1.177

Notes: The estimation method is OLS. Robust Standard errors in parentheses. Significant at *** 1%, ** 5% and * 10%. This table only reports the regressor of interest. Regressors not listed are those of Table 2. See main text for additional details.

patents belonging to USPC classes that were never designated as the primary class for patents receiving secrecy orders in our reference period. Panel B of Table 4 reveals that the treatment effect remains negative and statistically significant even within ‘non-secret’ patent classes.

In an unreported final test, we recompute the dependent variables by excluding forward citations originating from assignees that have received at least one secrecy order. These assignees are the most likely to be involved in sensitive technologies. The results remain quantitatively similar to the baseline findings.

In summary, these tests support the view that the decline in citations reflects a genuine reduction in the rate of invention. The results do not appear to be driven by inventors filing fewer patents while producing the same amount of inventions that are at risk of being subject to secrecy orders.¹⁹

5.5. Knowledge is in the air, but also in the patent document

Our previous analysis has shown that enforcing a secrecy order impedes the emergence of follow-on inventions, an effect that we interpret as being consistent with the disclosure effect in a broad sense. In the preceding section, we ruled out potential confounding factors and various alternative explanations for the observed effect. In this section, we delve deeper into the question of whether the disclosure of knowledge in the patent document plays a significant role in driving this effect.

Specifically, if the effect we observe is, in part, a consequence of restricted access to the codified knowledge contained within the patent document, we expect this effect to be more pronounced in areas geographically distant from the focal patent.

The literature on knowledge spillovers has long emphasized that knowledge is in the air, but that it is localized (see, e.g., Agrawal et al., 2008; Murata et al., 2013). The knowledge contributing to an invention is generally available for everyone to build upon. Consequently, closely related follow-on inventions are likely to emerge in geographically proximate entities, irrespective of the learning opportunities offered by the disclosure of codified knowledge through the patent system.

Therefore, to the extent that the filing of a patent facilitates the diffusion of valuable information, we anticipate that the effect of obstructing all channels of knowledge disclosure would be more pronounced for geographically distant pairs of patents.

To test this hypothesis, we first reconstruct our sample by exclusively focusing on focal patents owned by U.S.-based institutions. Subsequently, we construct five dependent variables, each counting the number of forward citations originating from assignees located within varying geographical radii: within 100 km, within 300 km, within 500 km, from 500 km to 1000 km, and more than 1000 km. Appendix Table B.2 presents the results of this set of estimates.

¹⁹ We also conducted an unreported robustness check to examine whether the negative effect of secrecy is more pronounced during periods of heightened geopolitical tensions, such as during the Cold War years. Our results showed no significant difference between the effects of secrecy orders imposed during the Cold War and those imposed afterward, suggesting that the adverse impact of secrecy is not limited to specific historical periods.

The substantial reduction in magnitude (and the loss of statistical significance) of the effect for closely located patent pairs aligns with the ‘knowledge is in the air’ hypothesis. Knowledge is highly localized, and it appears that assignees in relatively close proximity to the focal inventions do not critically depend on the patent document for learning.²⁰

The observation that the effect of secrecy primarily manifests in geographically distant patent pairs, where access to codified knowledge is more likely to be a significant channel of discovery, lends further support to the idea that information disclosure enabled by a patent is conducive to the diffusion of knowledge.

6. Conclusion

This paper shows that patents play a significant role in enabling knowledge disclosure. We offer robust evidence that secrecy orders hinder the development of knowledge related to classified inventions. Our analysis focuses on the citation patterns of 505 prior patents cited by secret patents, comparing them to a control group of 1,244 prior patents cited by non-secret patents. Our findings reveal a notable and statistically significant negative relationship between the enforcement of secrecy orders and the emergence of subsequent inventions.

According to our OLS estimates, patents cited by a patent subject to a secrecy order receive an average of 40 percent fewer forward citations than the control group when the secrecy order is in effect. Remarkably, this effect persists even after the lift of the secrecy order, suggesting a sustained impact resulting in a loss of inventive output. Furthermore, our main results remain robust across a wide range of alternative specifications, including the use of a text-based measure for tracking follow-on inventions. Unfortunately, we are unable to disentangle the precise sources contributing to this effect, whether it arises from the absence of published knowledge in the patent document or the suppression of peripheral disclosure channels.

These findings have important policy implications. Firstly, they underscore that the Invention Secrecy Act effectively achieves its intended outcome by significantly reducing the visibility of technological advancements associated with secret patents.

Secondly, our study contributes to the ongoing discourse on secrecy by shedding light on one of its social costs—the loss of subsequent inventions. Recent legislative changes in the United States, such as the expansion of prior user rights defense to patent infringement through the America Invents Act (AIA) (Khamin, 2014) and the adoption of the Uniform Trade Secrets Act by an increasing number of U.S. states (Seaman, 2015), have strengthened the incentives to keep inventions secret.

While our paper does not take a stance on whether current incentives lean too heavily towards secrecy, it highlights a tangible social cost of secrecy (see also Ganglmair and Reimers, 2019). Other costs, such as the duplication of R&D investments (Lück et al., 2020; Hegde et al., 2020), have been explored elsewhere. Conversely, scholars have identified benefits of secrecy, such as an increase in private R&D investments (Png, 2017). It is important to note that our analysis specifically examines ‘forced’ secrecy for a select set of technologies, and our results may not be directly applicable to situations involving voluntary secrecy.

In conclusion, the optimal level of secrecy to maximize overall growth remains an open question. This issue is far from settled, and further research is warranted to provide deeper insights into the complex interplay between secrecy, innovation, and growth.

CRedit authorship contribution statement

Gaétan de Rassenfosse: Conceptualization, Formal analysis, Methodology, Project administration, Writing – original draft, Writing – review & editing. **Gabriele Pellegrino:** Conceptualization, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Emilio Raiteri:** Conceptualization, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Data availability

Data will be made available on request.

²⁰ It should be noted that we do not construe this result as evidence suggesting that peripheral disclosure channels were not muted. While peripheral disclosure channels constitute one reason why knowledge is in the air, other factors also contribute. For example, actors in close geographical proximity may share a common knowledge base, tackle similar problems, and so forth.

Appendix A. Example of secrecy order documents

22

#8



**UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office**

Address : COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

Serial No. 631,044 Filed 07/16/84

Applicant DONALD L. SCOFIELD

Title VARIABLE MISSILE CONFIGURATION-STANDARD TELEMETRY SYSTEM **GROUP 220**

MAILED

JUN 24 1986

SECURITY ORDER

(Title 35, United States Code (1952), sections 181-188)

NOTICE: To the applicant above named, his heirs, and any and all his assignees, attorneys and agents, hereinafter designated principals.

You are hereby notified that your application as above identified has been found to contain subject matter, the unauthorized disclosure of which might be detrimental to the national security, and you are ordered in nowise to publish or disclose the invention or any material information with respect thereto, including hitherto unpublished details of the subject matter of said application, in any way to any person not cognizant of the invention prior to the date of the order, including any employee of the principals, but to keep the same secret except by written consent first obtained of the Commissioner of Patents and Trademarks, under the penalties of 35 U.S.C. (1952) 182, 186.

Any other application already filed or hereafter filed which contains any significant part of the subject matter of the above identified application falls within the scope of this order. If such other application does not stand under a secrecy order, it and the common subject matter should be brought to the attention of the Security Group, Licensing and Review, Patent and Trademark Office.

If, prior to the issuance of the secrecy order, any significant part of the subject matter has been revealed to any person, the principals shall promptly inform such person of the secrecy order and the penalties for improper disclosure. However, if such part of the subject matter was disclosed to any person in a foreign country or foreign national in the U.S., the principals shall not inform such person of the secrecy order, but instead shall promptly furnish to the Commissioner of Patents and Trademarks the following information to the extent not already furnished: date of disclosure; name and address of the discloser; identification of such part; and any authorization by a U.S. Government agency to export such part. If the subject matter is included in any foreign patent application, or patent this should be identified. The principals shall comply with any related instructions of the Commissioner.

This order should not be construed in any way to mean that the Government has adopted or contemplates adoption of the alleged invention disclosed in this application; nor is it any indication of the value of such invention.

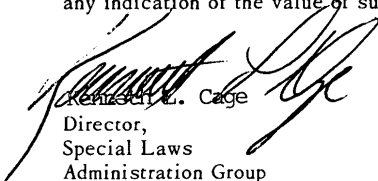

Kenneth L. Cage
Director,
Special Laws
Administration Group

Fig. A.1. Secrecy order imposition document.



**UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office**

Address : COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

SERIAL NUMBER	FILING DATE	FIRST NAMED APPLICANT	ATTORNEY DOCKET NO.
06/631,044	07/16/84	SCOFIELD	D 67592

ROBERT F. BEERS
OFFICE OF PATENT COUNSEL, CODE 012
NAVAL WEAPONS CENTER
CHINA LAKE, CA. 93555-6001

EXAMINER	
TARCZA, T	
ART UNIT	PAPER NUMBER
2202	

DATE MAILED: 07/23/92

RENEWAL OF SECRECY ORDER
[Title 35, United States Code (1952), Sections 181-188]

NOTICE: To the applicant(s), heirs of applicant(s), and any and all assignees, attorneys and agents, the designated principals.

The Armed Services Patent Advisory Board, Department of Defense (DOD), has notified the Commissioner of Patents and Trademarks that an affirmative determination has been made by a DOD agency, identified below, that the national interest requires renewal of the secrecy order. The secrecy order is therefore, renewed, effective for a period of ONE YEAR from the date of this renewal notice.

The secrecy order may be renewed for additional periods of not more than one year upon notice by a government agency that the national interest so requires.

DOD AGENCY: DARCOM NAVY AIR FORCE ASPAB

J. L. Bradshaw
Director,
Special Laws Administration

Fig. A.2. Secrecy order renewal document.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
06/631,044	07/16/1984	DONALD L. SCOFIELD	67592	7443

7590 07/14/2010
ROBERT F. BEERS
OFFICE OF PATENT COUNSEL, CODE 012
NAVAL WEAPONS CENTER
CHINA LAKE, CA 935556001

EXAMINER
TARCZA, THOMAS H
ART UNIT
PAPER NUMBER

3662

7-14-10

RESCINDING OF SECRECY ORDER
(35 U.S.C. 181-188)

The Secrecy Order, dated , prohibiting disclosure of publication of the subject matter of the above-identified application under 35 U.S.C. 181-188 is hereby rescinded. Normal prosecution is continued and any suspension thereof because of the secrecy order has been removed. This rescinding order does not affect the provisions of any classified government contract or existing laws relating to espionage and national security.

Note that this decision does not affect the removal of any national security markings. Applicant is required to obtain authorization to remove any such markings or pursue a new security order.

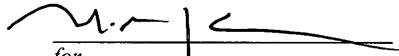

for
Director, Technology Center 3600
(571) 272-5150

Fig. A.3. Secrecy order rescinding document.

Appendix B. Secrecy orders by technology class

Table B.1
Distribution of patents with a secrecy order by technological class (top 25 USPC main classes).

USPC code	USPC definition	Freq.	Perc.
428	Stock material or miscellaneous articles	52	7.17
342	Communications: directive radio wave systems and devices (e.g., radar)	50	6.90
102	Ammunition and explosives	33	4.55
367	Communications, electrical: acoustic wave systems and devices	31	4.28
356	Optics: measuring and testing	26	3.59
250	Radiant energy	26	3.59
244	Aeronautics and astronautics	26	3.59
89	Ordnance	23	3.17
60	Power plants	22	3.03
427	Coating processes	21	2.9
372	Coherent light generators	21	2.9
264	Plastic and nonmetallic article shaping or treating: processes	18	2.48
364	Electrical Computers and Data Processing Systems	18	2.48
149	Explosive and thermic compositions or charges	17	2.34
333	Wave transmission lines and networks	17	2.34
359	Optical: systems and elements	16	2.21
343	Communications: radio wave antennas	16	2.21
501	Compositions: ceramic	13	1.79
156	Adhesive bonding and miscellaneous chemical manufacture	12	1.66
324	Electricity: measuring and testing	11	1.52
416	Fluid reaction surfaces (i.e., impellers)	9	1.24
423	Chemistry of inorganic compounds	9	1.24
429	Chemistry: electrical current producing apparatus, product, and process	8	1.10
228	Metal fusion bonding	8	1.10
340	Communications: electrical	7	0.97
–	Residuals USPC Classes (87 classes)	215	29.7
Total		725	100

Table B.2
Effect of secrecy orders on forward citations: counting citations by geographical distance.

	(1) < 100 km	(2) < 300 km	(3) < 500 km	(4) 500-1000 km	(5) > 1000 km
Panel A. OLS estimations					
Secret	-0.115 (0.131)	-0.090 (0.192)	-0.210 (0.215)	-0.466*** (0.125)	-0.774*** (0.244)
Panel B. Poisson estimations					
Secret	-0.023 (0.024)	-0.043 (0.073)	-0.111 (0.108)	-0.154*** (0.032)	-0.600*** (0.153)
Number of observations	1,033	1,033	1,033	1,033	1,033
Mean of dep. var.	0.546	0.909	1.201	0.684	2.323

Notes: Robust Standard errors in parentheses. Significant at *** 1%, ** 5% and * 10%. This table only reports the regressor of interest. Regressors not listed are those of Table 2. The dependent variable records the cumulative count of forward citations that the focal patent i has received from the year in which the secrecy order of its citing patent has been imposed until the publication year of the patent. We construct five dependent variables by counting citations within different geographical distances from the residence of the citing and cited assignees: within 100 km (col. 1), within 300 km (col. 2), within 500 km (col. 3), between 500 and 1000 km (col. 4), and above 1000 km (col. 5). Self-citations and foreign (not-U.S.) citations are excluded from the count of the dependent variables. See main text for additional details. Panel B: the Negative Binomial regression model did not converge; we have used Poisson instead. Marginal effects reported.

References

- Abadie, A., Spiess, J., 2022. Robust post-matching inference. *J. Am. Stat. Assoc.* 117, 983–995.
- Agrawal, A., Kapur, D., McHale, J., 2008. How do spatial and social proximity influence knowledge flows? Evidence from patent data. *J. Urban Econ.* 64, 258–269.
- Alcacer, J., Gittelman, M., 2006. Patent citations as a measure of knowledge flows: the influence of examiner citations. *Rev. Econ. Stat.* 88, 774–779.
- Arts, S., Cassiman, B., Gomez, J.C., 2018. Text matching to measure patent similarity. *Strateg. Manag. J.* 39, 62–84.
- Baruffaldi, S.H., Simeth, M., 2020. Patents and knowledge diffusion: the effect of early disclosure. *Res. Policy* 49, 103927.
- Belenzon, S., 2012. Cumulative innovation and market value: evidence from patent citations. *Econ. J.* 122, 265–285.
- Cox, G.W., 2019. Patent disclosure and England's early industrial revolution. *Eur. Rev. Econ. Hist.*
- de Rassenfosse, G., Raiteri, E., 2022. Technology protectionism and the patent system: strategic technologies in China. *J. Ind. Econ.* 70, 1–43.

- Dosi, G., Nelson, R.R., 2010. Technical change and industrial dynamics as evolutionary processes. In: Handbook of the Economics of Innovation, vol. 1, pp. 51–127. p. 00299.
- Furman, J.L., Nagler, M., Watzinger, M., 2021. Disclosure and subsequent innovation: evidence from the patent depository library program. *Am. Econ. J.: Econ. Policy* 13, 239–270.
- Galasso, A., Schankerman, M., 2014. Patents and cumulative innovation: causal evidence from the courts. *Q. J. Econ.* 130, 317–369.
- Ganglmair, B., Reimers, I., 2019. Visibility of technology and cumulative innovation: Evidence from trade secrets laws. ZEW-Centre for European Economic Research Discussion Paper.
- Graham, S., Hegde, D., 2015. Disclosing patents' secrets. *Science* 347, 236–237.
- Graham, S.J., Marco, A.C., Miller, R., 2018. The USPTO patent examination research dataset: a window on patent processing. *J. Econ. Manag. Strategy* 27, 554–578.
- Gross, D.P., 2022. The hidden costs of securing innovation: the manifold impacts of compulsory invention secrecy. *Manag. Sci.*
- Hall, B., Helmers, C., Rogers, M., Sena, V., 2014. The choice between formal and informal intellectual property: a review. *J. Econ. Lit.* 52, 375–423.
- Hegde, D., Herkenhoff, K., Zhu, C., 2020. Patent publication and innovation. SSRN WP No 3158031.
- Ho, D.E., Imai, K., King, G., Stuart, E.A., 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Polit. Anal.* 15, 199–236.
- Ho, D.E., Imai, K., King, G., Stuart, E.A., 2011. MatchIt: nonparametric preprocessing for parametric causal inference. *J. Stat. Softw.* 42, 1–28.
- Iacus, S.M., King, G., Porro, G., 2012. Causal inference without balance checking: coarsened exact matching. *Polit. Anal.* 20, 1–24.
- Iacus, S.M., King, G., Porro, G., 2019. A theory of statistical inference for matching methods in causal research. *Polit. Anal.* 27, 46–68.
- Iaria, A., Schwarz, C., Waldinger, F., 2018. Frontier knowledge and scientific production: evidence from the collapse of international science. *Q. J. Econ.* 133, 927–991.
- Jaffe, A., Trajtenberg, M., Henderson, R., 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Q. J. Econ.* 108, 577–598.
- Jaffe, A.B., de Rassenfosse, G., 2017. Patent citation data in social science research: overview and best practices. *J. Assoc. Inf. Sci. Technol.* 68, 1360–1374.
- Jones, B.F., 2009. The burden of knowledge and the “death of the renaissance man”: is innovation getting harder? *Rev. Econ. Stud.* 76, 283–317.
- Khamin, A., 2014. America invents act's prior user defense: lesson from global patent regimes and legislative history. *Pittsbg. J. Technol. Law Policy* 15, 132.
- Krugman, P., 1992. *Geography and Trade*. MIT Press.
- Kuhn, J.M., Younge, K.A., Marco, A.C., 2020. Patent citations reexamined. *Rand J. Econ.* 51, 109–132.
- Lemley, M.A., 2012. The myth of the sole inventor. *Mich. Law Rev.*, 709–760.
- Lück, S., Balsmeier, B., Seliger, F., Fleming, L., 2020. Early disclosure of invention and reduced duplication: an empirical test. *Manag. Sci.* 66, 2677–2685.
- Marshall, A., 1920. *Principles of Economics*. Macmillan.
- Moser, P., Ohmstedt, J., Rhode, P.W., 2017. Patent citations - an analysis of quality differences and citing practices in hybrid corn. *Manag. Sci.* 64, 1926–1940.
- Murata, Y., Nakajima, R., Okamoto, R., Tamura, R., 2013. Localized knowledge spillovers and patent citations: a distance-based approach. *Rev. Econ. Stat.* 96, 967–985.
- Nordhaus, W.D., 1967. The optimal life of a patent. Tech. Rep., Cowles Foundation for Research in Economics, Yale University.
- Ouellette, L.L., 2012. Do patents disclose useful information. *Harv. J. Law Technol.* 25, 545.
- Png, I.P., 2017. Law and innovation: evidence from state trade secrets laws. *Rev. Econ. Stat.* 99, 167–179.
- Rantanen, J., 2012. Peripheral disclosure. *Univ. Pittsbg. Law Rev.* 74, 1.
- Sampat, B.N., 2018. A Survey of Empirical Evidence on Patents and Innovation. NBER Working Paper, No 25383.
- Scotchmer, S., 1991. Standing on the shoulders of giants: cumulative research and the patent law. *J. Econ. Perspect.* 5, 29–41.
- Seaman, C.B., 2015. The case against federalizing trade secrecy. *Va. Law Rev.*, 317–394.
- Weston, J., Bengio, S., Usunier, N., 2011. Wsabie: scaling up to large vocabulary image annotation. In: Twenty-Second International Joint Conference on Artificial Intelligence.
- Whalen, R., Lungeanu, A., DeChurch, L., Contractor, N., 2020. Patent similarity data and innovation metrics. *J. Empir. Leg. Stud.* 17, 615–639.
- Younge, K.A., Kuhn, J.M., 2016. Patent-to-patent similarity: a vector space model. Available at SSRN 2709238.